VoIP essentials IP telephony veterans talk about five critical steps for ensuring the security of converged voice/data nets. **PAGE 9.**

VeriSign caves After weeks of criticism, the company agrees to suspend SiteFinder service. **PAGE 13.**

NetworkWorld

The leader in network knowledge ■ www.nwfusion.com   October 6, 2003 ■ Volume 20, Number 40

Security debate rages

Intrusion-detection critics and backers still sparring months after Gartner salvo.

■ BY ELLEN MESSMER

Strong aftershocks continue from the Gartner report that declared intrusion-detection systems dead and predicted the market for such products would be gone by 2005.

While the debate sparked by Gartner's assessment remains unresolved, reverberations are evident in the product road maps of IDS vendors. The companies are developing systems that can actively block attacks and passively detect them, a key recommendation in Gartner's report in June. Debates also are raging in corporate and government IT departments about whether to

> **❝It got a little ugly. Some IDS vendors said IPS vendors were bribing me.❞**
>
> **Richard Stiennon**
> **Vice president of research, Gartner**

buy IDS products.

Gartner's Vice President of Research Richard Stiennon stands behind his report's controversial conclusion — despite conceding a point or two to critics. And he remains surprised by the intensity of

the firestorm, which culminated in his being challenged in July before a collection of concerned federal agencies and unhappy IDS vendors.

"It got a little ugly," Stiennon says. "Some IDS vendors said [intru-

sion-prevention system] vendors were bribing me."

The "IDS is dead" report, as it's now widely called, stated IDS sensors used for passive monitoring of network traffic are a waste. According to Gartner, that's because they generate a lot of false alerts about attacks and are a round-the-clock management burden for IT. Declaring IDS a "market failure," the report advised Gartner clients to start blocking attacks outright instead of just monitoring for them, something the newer firewall-like devices — sometimes called intrusion-prevention systems (IPS) — can do. The number of IPS products is

---

Users banking on blades

■ BY JENNIFER MEARS AND DENISE DUBIE

Greater Baltimore Medical Center had a dilemma. Its IT needs were growing, but its data center space was not. After months of dead-end negotiations with vendors in an attempt to put multiple applications on

BladeServers
EARLY ADOPTERS

fewer, bigger boxes, the nonprofit turned in the other direction: It brought in blade servers.

"There were three business problems that really drove us toward the technology: One, we had no space in our data center, and we needed to add 30 servers. Two, we had [limited] power, and

---

Fortifying BGP: No quick fix

■ BY JIM DUFFY

In 1996 the U.S. government tapped BBN to develop a more secure version of the primary protocol used to route information around the Internet.

The effort was not in response to any particular data or network security breach. It stemmed from a realization that the Border Gateway Protocol (BGP) was becoming ever more vulnerable as the Internet grew in size and importance.

Yet seven years later, BBN's Secure BGP (S-BGP), which establishes a public-key infrastructure to stymie IP address spoofing, is still a work in progress and has yet to be implemented in

Diverging views

> **❝S-BGP is dead in the water.❞**
>
> **Fred Baker, Cisco Fellow and former IETF chair, and a proponent of soBGP, an S–BGP alternative**

> **❝Some of the options offered in soBGP would be disastrous from a security standpoint.❞**
>
> **Steve Kent, chief scientist for information security, BBN**

Introducing Microsoft Windows Server 2003. *Do more with less.*

You're being asked to do more. You're being asked to do it with less. Microsoft® Windows® Server 2003 is designed to help you manage these opposing forces with powerful server consolidation capabilities that increase efficiency, decrease man-hours, and lower your total cost of ownership. Download your free evaluation copy of Windows Server 2003 at microsoft.com/windowsserver2003 **Software for the Agile Business.**

Information Resources, Inc. (IRI) manages over 122 terabytes of data to provide consumer behavior insights, advanced analytics, and decision analysis tools for some of the largest consumer packaged goods, healthcare, retail, and financial companies in the world. To meet increasing demand for faster, more granular business intelligence while reducing costs, IRI is using 64-bit editions of Windows Server 2003 and SQL Server™ 2000 on an Intel Itanium 2 system to deliver faster answers to its customers. The result? IRI will be able to process more queries, using a fraction of the number of servers while realizing significant cost savings and improving customer service.

With 10 servers.

NetworkWorld

Page 57

Bradley Kuhn, left, and Chris Sontag, debate the GPL, page 45.

NetworkWorldFusion

www.nwfusion.com

Interactive

New standards page
The nice thing about standards is how many of them there are, right? Keep up with all the standards work with our new Standards page. **DocFinder: 7628**

News the way you want it
Track just the technologies, companies and authors you're interested in — we've now got more than 60 RSS feeds, including new ones on servers and viruses. **DocFinder: 7442**

Spam calculator
See how much spam is costing your company with our new tool. **DocFinder: 7930**

Seminars and events

Is your data center fully metered, fully accountable and fully effective?
The new data center is about ways to control and consolidate data, new ways to automate servers and storage, and new tools to maximize your network through virtualization. Discover them at The New Data Center: Powering the Enterprise, a Network World Technology Tour event. Free to qualified professionals. **DocFinder: 7534**

Columnists

Compendium
RSS for data-centric applications
RSS has taken off in the Weblog world due to its ease of use. Fusion Executive Editor Adam Gaffin introduces you to Jeremy Allaire, who's figured out a way to expand RSS into data-oriented applications. **DocFinder: 7931**

Telework Beat
One smart contingency plan
Net.Worker Editor Toni Kistner explains how the Aptela phone service helped an ISP maintain customer support during Hurricane Isabel. **DocFinder: 7932**

Small Business Tech
Small-business lovefest
Columnist James Gaskin says SMBs should take advantage of their size to get the best deals. **DocFinder: 7933**

Home Base
Virtual company, rock-solid assistance
Columnist John Brandon details how one professional relies on Groove P2P collaboration to manage her many clients. **DocFinder: 7934**

Breaking News

Come online for breaking news every day. **DocFinder: 6342**

Free e-mail newsletters
Sign up for any of more than 40 newsletters on key network topics. **DocFinder: 6343**

What is DocFinder?
We've made it easy to access articles and resources online. Simply enter the four-digit DocFinder number in the search box on the home page, and you'll jump directly to the requested information.

ALLEN CRAWFORD/PLANKTON ART CO

New Bits

Microsoft security practices face legal action

■ A California lawyer is trying to organize a class-action lawsuit against Microsoft, saying the company is engaging in unfair business practices in violation of state law because it has failed to secure its software against worms and viruses. The suit was filed in Los Angeles on behalf of a woman who is seeking undisclosed damages and an order requiring Microsoft to improve its security notification system, which has been a constant target of criticism from end users. The woman says her name and Social Security number were stolen and used fraudulently by someone who hacked into her computer. Dana Taschner, a Newport Beach, Calif., lawyer who filed the suit, wants to expand it into a class-action case against Microsoft. Traditionally, software makers have been shielded from such legal actions because of licensing agreements that software users sign. Microsoft said it is reviewing the complaint.

Deadline passes on H-1B visas

■ A congressional cap on the number of foreign workers allowed into the U.S. on H-1B visas has fallen to pre-dot-com boom levels after Congress failed to reauthorize a higher limit before last week's deadline. But debate on the program is not over. Congress still could act to increase the number of H-1B visas even though the 2004 fiscal year has started with a cap of 65,000, some supporters and opponents of a higher cap say. Intel will continue to press for a higher cap and to have advanced-degreed engineers exempted from the cap, says Tracy Koon, a company spokeswoman. "It's clear when you look at U.S. graduation numbers, there's a shortage there." Even opponents of a higher H-1B cap acknowledge the fight isn't over although the congressional deadline passed. H-1B critic Norm Matloff, a computer science professor at the University of California at Davis, says he expects the issue to resurface early next year.

What IT will look like in 2010

■ The average IT department will look radically different by 2010, full of relationship managers and "touchpoints" between the company and its outsourced suppliers, Gartner Vice President Ian Marriott said at a roundtable last week in London. He said the trend toward outsourcing, including outsourcing business processes and more standard IT services, would involve a complete change of mindset for the IT manager. "The IT department will need fewer technical skills and more business skills," he said. Persuading management of the need for outsourcing is going to be hard because a huge investment in staff is needed. "You won't be able to retrain everyone in the IT department, and so you'll need board commitment to spend the money upfront. You should be spending at least 5% of the value of the outsourcing deal on just managing that deal, or it just won't be good enough," he said.

SCO turns legal guns on SGI

■ Silicon Graphics Inc. is the latest technology company to be dragged into The SCO Group's dispute over the Linux operating system. In an Aug. 13 letter addressed to SGI's legal department and released to the media last week, SCO CEO Darl McBride said SGI's contributions to Linux put it in breach of its 1986 Unix licensing agreement, originally signed with AT&T but subsequently transferred to SCO. "SGI flagrantly permitted the copying and use of our proprietary information without any knowledge of the identities of the

TheGoodTheBadTheUgly

😇 **Kissing and making up.** Cisco and Huawei Technologies, the big 3Com partner based in China, last week made an agreement that could soon end the lawsuit between the companies. They agreed to stay litigation in a suit Cisco brought against Huawei, saying the latter pirated Cisco IOS router code and lifted material from Cisco owner's manuals.

👿

HP CEO Carly Fiorina for the sixth consecutive year has topped Fortune's list of most powerful businesswomen. No offense intended to Fiorina, but doesn't the fact that any woman would be chosen for this distinction six years running speak to the general dearth of women in executive suites as much as it does the prowess of that particular executive?

BRIAN GAIDRY

🌐 **Under attack.** The rate of network-based attacks rose 19% in the first half of the year vs. the first half of last year, according to a new report from Symantec. Of great concern, the company says, is the rise in blended threats, which combine multiple break-in techniques and accounted for 60% of malicious code submissions in the first half of this year.

recipients" and "subjected our source code to unrestricted disclosure, unauthorized transfer and disposition, and unauthorized use and copying," McBride said in the letter. The letter threatens to terminate SGI's Unix license as of Oct. 14 should SGI fail to "remedy all violations." SGI responded to SCO's letter in early September with a letter saying SCO's allegations were without merit, SGI spokeswoman Marty Coleman says.

Storage organizations debut

■ Two user groups for storage professionals debuted last week. The Association of Storage Networking Professionals, sponsored by a for-profit company, and StorageNetworking.Org, headed by the vendor-neutral Storage Networking Industry Association and the Information Storage Industry Center at the University of California, San Diego will provide members with educational opportunities and the information to make knowledgeable storage acquisitions. The ASNP will have membership rates of $190 per year; StorageNetworking.Org will be free.

Five tips for securing a converged net

■ BY PHIL HOCHMUTH AND TIM GREENE

IP telephony and voice over IP are by no means the standard for carrying enterprise voice just yet. But these technologies have been in the real world long enough for users to have learned some tricks for protecting a converged infrastructure against network threats, both external and internal from inside the firewall.

From their work in the field, these IP telephony veterans — including users and consultants — give the following five essential tips for ensuring security in a converged voice/data infrastructure.

1. Start with the basics.

"Viruses are a paramount concern in any organization, on any server or application," says Ray Ortega,

senior consultant for voice at ThruPoint, an integrator of IP telephony networks. "One key thing users need to do is to stay on top of their [IP PBXs] and make sure they have the latest virus protection and patches applied. One thing we make sure of is that clients know that call servers need to be maintained and monitored."

Ortega says common-sense precautions such as intrusion-detection software and good firewalls go a long way in keeping IP voice running in the event of a network or virus attack. Disabling or limiting Web access to phones and IP PBXs — as many of these devices run mini-Web servers for management purposes — is another step.

CERT, the independent network security organization, recommends filtering inbound traffic to Session Initiation Protocol (SIP) devices and denying traffic to those

devices that are not intended to handle public services. Similarly, such boxes rarely need to initiate SIP sessions, so filtering outbound traffic that is initiating sessions can prevent these machines from being used to launch attacks, CERT says.

2. Treat phones as IP clients.

Intruders can spoof IP addresses to make illicit gear seem to be a trusted device, which can then intercept traffic. This behavior can be blocked by requiring endpoint IP address authentication, many users and experts say.

Businesses want to make sure users authenticate themselves to the network before using IP phones, says Iain Stevenson, service director for access at analyst firm Ovum. "With VoIP you have a highly portable client [softphone]. Authorizing the user is very important so you don't get anyone dialing out and making long, international calls."

Many experts also recommend setting up logon prompts and PINs for IP phones. Disabling auto-configuration of IP phones from an IP PBX could also be useful in preventing unauthorized IP telephony clients from making calls via your IP PBX.

3. Keep converged voice and data separate.

It sounds contrary, but completely mixing voice and data can be bad, some experts say.

"You want to limit the kind of communication on IP PBXs strictly to the devices that they need to communicate with," Ortega says. These include IP phones, PCs running softphones and application servers that need to interact with the IP PBX, such as mail or contact center servers.

He also recommends putting voice on its own virtual LAN segment, and always running site-to-site VoIP links over dedicated WAN circuits.

This is the practice of one aerospace parts manufacturing company, with offices on the East Coast and in Europe. The company uses IP PBXs in its branch offices, which are connected by private ATM links leased from a carrier. "We don't let any [VoIP] traffic go beyond our private LAN and WAN," an IT administrator for the firm says.

4. QoS as a security measure.

While many VoIP users recommend quality of service (QoS) to maintain voice quality, making voice packets a high priority can help in the event of a security incident, some say.

"We use Layer 3 switches that give voice the highest priority over any other traffic types," says John Orbaugh, director of MIS for the Tyler Independent School District in Texas. The school deployed Nortel's Business Communication Manager — a branch IP PBX — to four high school campuses, with a Nortel Meridian PBX acting as a central call switch.

Orbaugh says he uses network security basics — firewalls and some intrusion detection.

"With QoS we should still be able to push voice through" in case of a network-saturating worm or denial-of-service attack, he says. "Quality might go down a bit, but it would still be up and running."

5. Encrypt calls where it makes sense.

Eavesdropping by people intercepting the voice stream is possible and

Brightmail's anti-fraud services target spoofers

■ BY CARA GARRETSON

Catering to corporate customers who want more than just junk-mail protection from their anti-spam vendor, Brightmail last week announced a new anti-fraud service and the addition of Symantec's anti-virus software to its gateway product.

Brightmail's anti-fraud service aims to minimize the damage caused by fraudulent e-mail campaigns that "spoof" a company's brand. Spammers launch e-mail campaigns that look like messages sent by well-known companies, realizing that recipients are more likely to open an e-mail from an established brand than from a random name.

Sometimes these e-mails will send recipients to a Web site also branded with the company's logo, ask them to enter personal information, including Social Security number, and steal their identity, says Paul Bruno, product manager at Brightmail. The company estimates approximately 10% of the e-mail scanned by its Probe network — a collection of 2 million decoy e-mail accounts the company uses to gather intelligence on spam — are fraudulent.

The company's new service detects e-mail fraud, usually by figuring out that the Web link embedded in the e-mail links to a phony Web site. Brightmail then alerts the company whose brand has been stolen and blocks the fraudulent messages from arriving in mailboxes of its anti-spam software users.

Brightmail is in a good position to offer anti-fraud services because its Probe network captures a large quantity of unsolicited e-mail, one analyst says. "They're watching what's going on, and they're in a

position to be able to track [fraudulent e-mails] and provide this service," says Jan Sundgren, an industry analyst at Forrester Research.

Brightmail also announced last week an upgrade to its Anti-Spam Enterprise Edition software that includes a version of Symantec's gateway anti-virus offering. The anti-virus software is sold as a separate module, says Carlin Wiegner, Brightmail's enterprise product manager.

The company also has increased the software's ability to catch spam in Version 5.1 with the addition of heuristic technology, Wiegner says. Heuristics identify unwanted e-mail messages by looking for telltale signs common in spam, such as excessive use of capital letters or multiple colors.

The addition of heuristics makes Brightmail's software more effective in catching spam without adding to the number of false positives, or wanted e-mail, that the software catches. "We're in a position now where we feel we've invented better [heuristic] technology to get us to 95% accuracy," Wiegner says. With Version 5.0, Brightmail's accuracy rate for catching spam was 90%.

Brightmail's enterprise software, which began as a product for ISPs, competes with packages from companies such as Active State, Cloudmark and Sunbelt Software, and with services from Front-Bridge Technologies, Postini and others.

Brightmail's Anti-Spam Enterprise 5.1 software costs about $1,500 per year for 49 users; additional users cost $14 per user, per year. Pricing for the Symantec anti-virus module will be announced next week. No pricing was available for the anti-fraud service. ■

Polycom unit boasts better video compression

■ BY JASON MESERVE

Polycom this week will introduce its first videoconferencing appliance that supports a new video compression standard, which is said to cut in half the bandwidth required to provide the same quality video as the previous standard.

With H.264 built-in, the VSX 7000 entry-level system lets Polycom catch up with rival Tandberg, which offers H.264 in all its video endpoints. The previous standard, dubbed H.263, also is supported in the VSX 7000.

"When we start doing international calls [with H.264], we'll see the benefits," says Alex Nason, senior manager of business development at Johns Hopkins International, an affiliate of Johns Hopkins medical center in Baltimore. Nason says some overseas calls can cost $12 per minute for a 384K bit/sec connection using six bonded ISDN lines. "If I can reduce [by] two lines [a $4 per minute savings], double the quality and still save money, that's a no-brainer."

Nason, who's tested the VSX 7000 for a few weeks, also likes the new user interface on the device. The interface features a stream-

Polycom's new VSX 7000 supports the new H.264 video-compression standard that delivers the same quality video as previous codecs at half the bandwidth.

lined main menu and reconfigured remote control, which makes it more intuitive for end users to make a call. An administrator can further simplify the interface by eliminating or adding options. One seemingly minor addition to the remote control — a "dot" button — makes it much easier to key in IP addresses, Nason says.

In a departure from its traditional appliance setup, where audio typically is delivered from the speakers on the connected TV, Polycom has added a built-in speaker and a detached subwoofer to the VSX 7000 and support for its Siren14 audio protocol. Siren14 delivers audio at 14 KHz, double the range of other audio protocols. Other features include dual-monitor emulation for simulating two monitors on one large screen and support for 11 languages. Administrators can access the Web-based management system on the device in one language, while end users see all the navigation screens in their native language.

The VSX 7000 will begin to replace Polycom's older ViewStation MP, SP and 512 lines although the company still will make and support those units, says Maggie Smith, director of product management.

The base VSX 7000 costs about $6,000, runs IP (H.323) only and comes with a 360-degree microphone, subwoofer and remote. Users can add ISDN (H.320) support via an optional hardware module for about $2,000. Multipoint support for connecting up to four end users in a conference with voice-activated presence (who's speaking is shown on screen), costs about $3,000. The Visual Concert hardware option for connecting a laptop and sharing data in a conference costs about $1,500. ■

BGP

Internet routers. Router memory constraints, processing overhead concerns and the downtrodden state of the telecom economy are cited as reasons why.

"The state of security in BGP is pretty minimal," says Alex Zinin, area director of the routing and sub-IP working groups in the Internet Engineering Task Force (IETF). "As it is deployed today, there is no mechanism to authenticate and identify the authorization of a specific [routing information] announcement."

What's more, work on BGP security is more divided than united (see related story, below). Cisco and some ISPs are working on an alternative to BBN's S-BGP, called Secure Origin BGP (soBGP), which authenticates yet also lets ISPs implement routing policy.
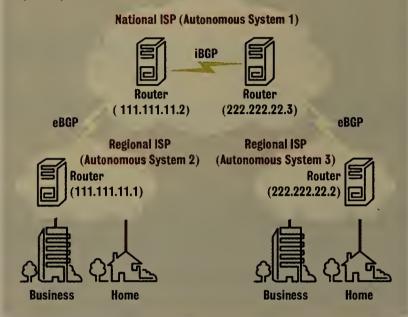
"S-BGP is dead in the water," says Cisco Fellow Fred Baker, former chair of the IETF.

That's an assertion to which Steve Kent, BBN's chief scientist for information security, counters: "Some of the options offered in soBGP would be disastrous from a security standpoint. There are concerns that soBGP doesn't architecturally nail things down."

Security isn't the only concern with BGP. Other public and private efforts have sprung up to address BGP's perceived shortcomings in scalability and reliability as traffic on the Internet

How BGP works

BGP serves as the Internet's primary routing protocol, performing interdomain routing within and between autonomous systems, which are networks or groups of networks under a common administration with common routing policies. BGP peer or neighbor routers initially exchange their full BGP routing tables, including autonomous system numbers, IP addresses and routes. Routers exchanging BGP updates within an autonomous system run internal BGP (iBGP) sessions, whereas those in different autonomous systems run external BGP (eBGP) sessions.

National ISP (Autonomous System 1)

iBGP

Router
(111.111.11.2)

Router
(222.222.22.3)

eBGP

eBGP

**Regional ISP
(Autonomous System 2)**
Router
(111.111.11.1)

**Regional ISP
(Autonomous System 3)**
Router
(222.222.22.2)

Business   Home   Business   Home

continues to double each year.

Some say it's time to move beyond the 14-year-old protocol, while others say doing so would cause drastic disruptions to the thousands of routers in and providing access to the Internet.

"A whole new protocol tends to make people think significant investment and high risk," says Martin Capurro, senior director of product management at Qwest. "We'd like to see a solution that just enhances the current one."

Proposed enhancements are plentiful. For reliability, the IETF and a number of router vendors developed so-called non-stop

routing/forwarding and graceful restart extensions to BGP to keep data flowing as the protocol resets (see www.nwfusion.com, DocFinder: 7941).

But ISPs are highly selective when it comes to incorporating such revisions.

Packet Design, a start-up led by industry veteran Judy Estrin, learned this firsthand.

The company last year unveiled BGP Scalable Transport, a protocol for streamlining communication of BGP routing information. By reducing the number of TCP connections required between routers, the technology could improve scalability and lessen security risks and the effect of lost connections, Estrin says.

But this Packet Design technology never caught on.

"We felt that the routing vendors just did not seem to want to spend the energy on fixing BGP," Estrin says. "The service providers were in enough disarray in terms of reorganizing and consolidation, [and] they didn't feel that they could put significant pressure on the routing vendors to get the capability. We couldn't deploy it without a router."

And router vendors found no need for such a technology.

"It's not something that we, as an implementor of the protocol, ever felt necessary to avail ourselves of," says Matthew Kolon, senior solutions engineer at Juniper.

As a general-purpose protocol,

Alternatives address BGP problems, but do they add their own?

The two proposals for addressing Border Gateway Protocol's security shortcomings might have some of their own.

BBN's Secure BGP (S-BGP) is intended to address a "fundamental problem" with BGP: the authenticity of routing update information, according to Steve Kent, BBN chief scientist for information security.

"What makes security for BGP tricky is that generally, this update information is transitive," he says. "One ISP is saying to another, 'I received this routing information from one of my neighbors with regard to this chunk of address space. If you want to send traffic for this chunk of address space to me, this is the path it would take.' Today, there's just no security for that. There's no way for the receiver to tell whether the update that's received is authentic."

S-BGP seeks to establish a public-key infrastructure that uses digital certificates to authenticate two pieces of data: which chunks of address space have been allocated to them and what autonomous system numbers have been allocated to them.

But S-BGP inhibits an ISP's ability to establish policy for its routers, says Cisco Fellow Fred Baker, whose company, along with ISPs such as Genuity, have written an alternative called Secure Origin BGP (soBGP).

"[With S-BGP, the] downstream service provider cannot apply a policy that says, 'I'm going to accept this prefix from you but not that one,'" Baker says. "It fundamentally breaks BGP's ability to be used in a policy system where you might redivide the information. S-BGP is the right concept, but it's put together in a way that an ISP can't really effectively use."

The soBGP proposal is an effort to let ISPs authenticate route advertisements and implement policy on them. But according to Kent, soBGP provides too many ways to do certain things, which when implemented differently, hamper interoperability.

The Internet Engineering Task Force is acting as mediator in the S-BGP/soBGP dispute. The routing protocol security working group within the IETF's Routing area is developing a so-called threat model that attempts to document the security requirements for Internet routing systems.

This work might provide the middle ground on which S-BGP and soBGP backers can come to a resolution, says Alex Zinin, director of the IETF's Routing and Sub-IP areas.

"This should help bring people on the same page as far as what we actually expect from routing protocols from a security perspective, as opposed to each protocol designer or each service provider making their own conclusions and assumptions," he says.

— Jim Duffy

VeriSign suspends controversial service

■ BY JORIS EVERS

Having stood firm for weeks under a barrage of criticism, VeriSign last Friday agreed to suspend its controversial Site-Finder service after the Internet's primary governing body issued an ultimatum that it do so or face legal action.

"We will accede to the request while we explore all of our options," VeriSign spokesman Tom Galvin told Reuters News Service.

Earlier in the day, the Internet Corporation for Assigned Names and Numbers (ICANN) informed VeriSign in writing that it had until 9 p.m. Oct. 4 to comply. In that letter, ICANN Chief Executive Paul Twomey says the changes implemented by VeriSign in SiteFinder "have had a substantial adverse effect . . . on the stability of the Internet."

ICANN's Security and Stability Advisory Committee is calling for public comments and reports on SiteFinder and has scheduled a "fact-gathering meeting" for Oct. 7 in Washington, D.C. The committee plans later to issue a report on the effects of SiteFinder on the stability of the Internet.

"VeriSign introduced its wild-card service, and although I am told there was some kind of advance notice, it came upon the world as a surprise, and after some hours it became clear that it is a pretty big deal," says Steve Crocker, chairman of ICANN's Security and Stability Advisory Committee. "The preliminary evidence suggests SiteFinder has impacted the stability of the Internet."

VeriSign controls the main database of .com and .net domain names (see related story, page 14). Last month the company added a wild card to the databases, sending Web users who enter a nonexistent .com or .net address to SiteFinder, a new service that offers Web links and paid advertisements.

SiteFinder has drawn a storm of criticism from technical and commercial fronts. At least two competing Internet companies have sued VeriSign, charging unfair competition. Some anti-spam filters failed when SiteFinder was used, and an uproar among network administrators prompted the Internet Software Consortium to update its DNS software so SiteFinder could be blocked. ISPs are said to have made changes to their networks to bypass the VeriSign service.

"Now you have a warring set of changes and it becomes a rickety system. That makes us engineers nervous. Generally we like to make changes slowly, carefully and with a great deal of consultation," Crocker says.

VeriSign had said prior to Friday's concession that it supported discussion on SiteFinder.

"We certainly are in favor of the community having a healthy discourse on all of the technologies and innovations on the Internet; SiteFinder is one of those. We are looking at how the [Oct. 7] meeting is shaping up and we will make our determination on participation based on that," says VeriSign spokesman Brian O'Shaughnessy.

Evers is a correspondent with the IDG News Services' San Francisco bureau.

BGP

BGP contains the features necessary to implement the scalability and security features appropriate between ISPs, Kolon says.

"A lot of it has to do with the implementation," he says. "[Limitations are] related not to the protocol itself but to the business and political relationships that are inherent in interdomain situations."

Another vendor echoes Kolon's views. Proficient Networks makes network appliances and software designed to reduce routed WAN infrastructure costs and improve application performance so service providers can deliver on service-level agreements (SLA).

"BGP really is about implementation — not necessarily flaws with the protocol," says Allan Leinwand, Proficient co-founder and president. "There are definitely some security hooks in BGP, such as MD5 checksums or digests on the information, but no one seems to be using them. So there's definitely a way for BGP peers to authenticate each other and to verify that the data coming across is valid and not being hacked or spoofed or replayed."
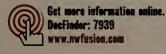
Service providers probably are not implementing features such as MD5 because they are not included in years-old BGP templates used to establish peering and interdomain policies, Leinwand says. ISPs also might be concerned that MD5, which requires additional CPU horsepower on the router, could sap performance and potentially violate customer SLAs, he adds.

"Do you want to add features and functions that are more CPU-intensive to BGP if you're already a little worried about the fact that it's perhaps not scalable?" Leinwand asks.

MCI uses MD5 to get a little more out of BGP. The carrier employs it on a per-request peer basis in some cases, says Jennifer Brooks, director of global IP engineering at the carrier's ISP operations.

"The clear thing to explore is, what is the risk you assume will occur with compromising BGP or hijacking a BGP connection?" she says. "[Hijacking is] a very hard thing to do. It's not something you can do over a standard BGP connection between two peers. From a hacker perspective, it's difficult unless a lot of information is provided."

AT&T devised its own method for dealing with BGP route-table integrity. In addition to route filtering, it developed a peering monitor that inspects information sent to its network from other parts of the Internet.

PeerMon, as AT&T calls it, looks for cases in which others are misrepresenting the carrier's address block. AT&T then can notify the unsuspecting ISP that it might need to reconfigure its network; or if the misrepresentation is of a malicious origin, attempt to track down the perpetrator.

"There's no authoritative list of who owns what address block," says Jennifer Rexford, a technology consultant at AT&T Bell Labs. "So when a piece of information is sent into the protocol, all it really takes is someone typing incorrectly or intentionally typing incorrectly to put misinformation into the protocol. Even if BGP could check a very accurate repository of that information, it might be extremely slow."

Qwest uses MD5 in all internal BGP sessions with its peers and to authenticate MPLS Resource Reservation Protocol connections, Capurro says. Qwest also separates BGP route reflection functions from the packet forwarding router onto separate servers, he says.

That way, the carrier also can separate public and private traffic, which frees up CPU cycles, increases memory and minimizes security risks, Capurro says.

BGP is adequate for the interdomain routing infrastructure of the Internet, but a new protocol is needed to swap control information for IP- and MPLS-based services such as VPNs, MCI's Brooks says.

"There's a lot of controversy right now in the IETF about the scalability of BGP," she says. "The requirements of the services for new features and enhancements are impacting the overall BGP source code. Should we cap BGP where it is today and create a new protocol that would be used for [RFC] 2547 [VPNs]?"

BGP wasn't originally intended to support VPNs, but was extended to accommo-date them, Brooks says. Some of these extensions could affect the protocol's performance and interoperability, she says.

"You're always going to see that collateral damage as you keep the two [service and infrastructure functions] together," Brooks says. "We're loading BGP down with more and more features than it really was intended to use."

While backing a new protocol for services, Brooks says BGP should remain the operational protocol exchanging information between autonomous systems in the Internet.

"I don't see that changing," she says. "It is too nested into the networks themselves to ever be able to safely undo that. There would be major routing instability if you were to try to move away from BGP."

The IETF doesn't see things changing anytime soon either.

"At this point, we're not defining a new routing protocol," IETF's Zinin says. "And we're not actively working on a new Internet routing and addressing system." ■

Get more information online.
DocFinder: 7939
www.nwfusion.com

More online!

The other VeriSign domain story. Discuss VeriSign's "wildcarding" policy in our forum.
DocFinder: 7999

Start-up talks up voice for WLANs

■ BY JOHN COX

Stop thinking about wireless LANs for data and start thinking about them for wireless voice.

That's the pitch from the latest WLAN start-up, Meru Networks, which this week is shipping an access point and controller intended to make wireless voice-over-IP practical for corporations.

The devices run Meru's Air Traffic Controller, which includes algorithms designed to boost the efficiency of any 802.11 WLAN and optimize it for voice traffic. The software does this by exploiting some lesser-known or less-used features of the 802.11 protocol, says Kamal Anand, vice president for marketing and sales.

"There is not a strong business case today for having WLANs throughout an enterprise," he says. "But voice [over wireless] will make that case."

Anand says Air Traffic Controller does three things for a WLAN:

● Boosts the number of users who can use an access point without performance loss from less than 20 to nearly 100.

● Automatically detects voice calls, makes a set of adjustments for them and gives voice traffic priority on the radio links.

● Slashes the time needed to hand off an end user, or a call, from one access point to another from around 300 milliseconds to nearly zero.

The 802.11 protocol accesses the radio wave via a contention technique: if one client detects another, it will back off and try again. Meru's algorithms make this contention a highly systematic process, orchestrating connectivity for scores of WLAN clients to a single access point.

So instead of all the clients "speaking at once," the software lets each one speak in turn, Anand says. This results in less delay and improved performance.

The software prioritizes voice traffic by detecting a voice call via Session Initiation Protocol, and then allocates bandwidth accordingly.

The company's software streamlines handoffs by grouping multiple Meru access points into one "virtual" access point with one Basic Service Set Identifier, which is the media-access-control address. The Meru algorithms let its software jump ahead to the next radio device, anticipating the moving WLAN client. The software sidesteps the need to completely set up and tear down handoffs between each separate access point.

The Meru controller works like the emerging class of WLAN switches rolled out in recent months by a slew of other companies, which centralizes authentication, encryption, management and rogue access-point detection. The controller plugs in via Gigabit Ethernet ports to a core switch; a discovery protocol lets the Meru access points find their way to the controller over the existing corporate network.

Access points have a list price of $595; controllers cost $8,000. ■

PROFILE:
MERU NETWORKS

Location:	Sunnyvale, Calif.
Founded:	February 2002
Primary product:	WLAN access point and controller, with proprietary algorithms for optimizing 802.11 networks for voice traffic.
CEO:	Ujjal Kohli, chairman and founder; formerly head of sales and marketing for AirTouch Cellular.
Financing:	Funded by seven venture funds, including Clearstone Venture Partners, NeoCarta and Evercore Ventures; amount not disclosed.
Competitors:	WLAN switch vendors such as Aruba, Airespace, Proxim, Trapeze; and Cisco, Extreme, and other network hardware vendors.
Fun fact:	In the mythology of India, "Meru" is the name of the sacred mountain that is the center of the universe; it represents aspiration.

Easing global domain name use

■ BY CAROLYN DUFFY MARSAN

The biggest barrier to widespread corporate use of internationalized domain names is the lack of support in key applications such as Web browsers and e-mail clients. So VeriSign, the central registry for domain names ending in .com and .net, is leading an attempt to lower that barrier.

VeriSign will announce as early as this week the formation of the IDN Software Consortium, which will promote the development of IDN-compliant software.

Today, Internet users have to download special plug-ins to resolve IDNs. VeriSign's goal is for software developers such as Apple, Lotus and Microsoft to provide built-in IDN support so that plug-ins are no longer required.

"The [IDN] standards are great, but we have to get the software developers and applications to support them," says Ben Turner, vice president of naming services at VeriSign.

VeriSign officials would not identify the companies that have joined the consortium but said the group will have its first meeting next month.

With IDNs, multinational corporations can create native-language Web sites to market products in each country where they conduct business. Karlsburg Brewing in Germany and Coca-Cola's Korean subsidiary both use native language domain names.

The IDN Software Consortium is good news for advocates of IDNs, who have labored for several years to develop a scheme for processing multilingual domain names without disrupting the Internet's DNS.

"It's excellent," says Paul Hoffman, one of the authors of the IDN standards, which awaits final approval from the Internet Engineering Task Force (IETF). "VeriSign has already made their IDN tool kit freely available to developers. With these tools being available and pushed by the consortium, developers will say that it's not so hard for them to support the standard."

Protecting brand names

Until now, domain names and e-mail addresses were based on English language, ASCII characters. Internet users overseas are clamoring to surf the Web and send e-mail using their native languages rather than the English approximations used today.

Many U.S. companies bought IDNs to protect their brand names and trademarks in other languages. However, few of those IDNs link to actual Web sites because of the technical difficulties involved.

"Overall, the number of domain names that have been sold is about 50 million worldwide," says Jay Westerdal, president of consulting firm Name Intelligence. "About 5 million of those names were in foreign languages. Most of the IDNs that were sold are in Asian languages such as Korean, Japanese and Chinese."

Westerdal says U.S. companies aren't using IDNs because of the need for software plug-ins.

"I was in China, and every time I saw a Web site

Looking abroad

The IETF's three specifications for resolving IDNs are:

RFC 3490
Internationalized Domain Names in Applications (IDNA), which describes a standard mechanism for handling non-ASCII characters in the DNS.

RFC 3491
Nameprep, which specifies rules for processing IDNs using the Unicode standard.

RFC 3492
Punycode, which encodes a Unicode string into an ASCII string.

being advertised, it was in English," Westerdal says. "The fact that users need a software plug-in is a real hurdle for lots of companies."

The IDN Software Consortium wants to remove that hurdle by helping software developers comply with the IETF's specifications for resolving IDNs. The IETF has developed three protocols that convert foreign language characters into Unicode, a computer industry standard, and then encode them in ASCII for transmission over the Internet's DNS.

IDN generates little backing

Since their release in March, the IETF's IDN protocols have generated little support among developers. Released this summer, Netscape 7.1 (also known as Mozilla 1.4) is the only Web browser with built-in support for IDNs. Microsoft will not comment on plans to support IDNs in Internet Explorer, Outlook or Outlook Express, and Lotus says it is not working on IDN support for Notes.

The IETF's IDN standards are solid, according to Hoffman, who last month ran an interoperability test of eight internationalized software packages from all over the world. The software packages — including Web browsers and plug-ins, e-mail clients, zone editing programs and programming tool kits — were run though 120 tests. Hoffman expects IDN-compliant software to ship by year-end.

Until then, domain name registries and registrars are forging ahead with sales of IDNs. More than 40 registrars offer domain names in 350 local languages for .com and .net.

VeriSign offers Internet users two ways to get to these IDNs: a free software plug-in called i-Nav or a Web navigation service, both of which support the IETF's specifications.

"People have downloaded 12 million plug-ins since January," Turner says. "Companies in Korea, China and Japan are just starting to use IDNs. We're seeing several hundred thousand [IDN] resolutions a day."

Turner estimates that more than 50,000 .com and .net IDNs point to multi-page Web sites.

Meanwhile, Afilias, the registry for .info names, announced last month the availability of standards-compliant IDNs using German script characters. Afilias has sold 260,000 .info names in Germany, which has emerged as the No. 2 market for .info names.

"The world is multilingual and multicultural," says Ram Mohan, vice president of operations and CTO of Afilias. "Businesses and applications on the Internet must go that way or risk becoming irrelevant." ■

Correction
■ The story "Raising an RFID ruckus" (Sept. 29, page 73) should have spelled a venture capitalist's name as Bob Hower.
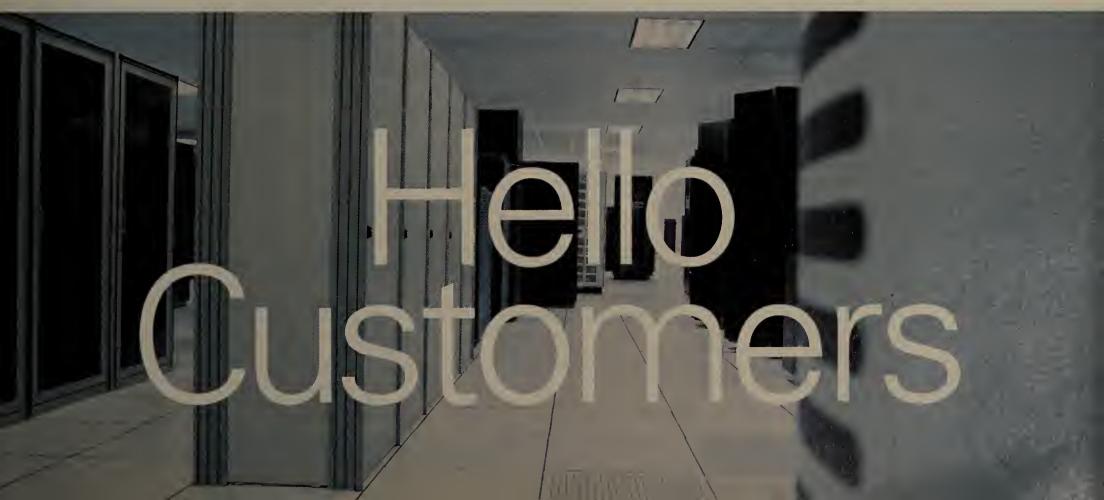
Goodbye Hackers

The right management should do more than just protect. It should also enable.

eTrust™ Security Management Software

In the world of on-demand computing, it's vital that your IT environment be both secure and accessible. That's why it's essential that you have the right security management software. With eTrust security management software, you get the very best in access, identity, and threat management all seamlessly integrated with your existing technology. On the one hand, you can rest assured, knowing that your information is safe from prying eyes. At the same time, you don't have to worry about partners, customers or employees being locked out of areas that they need to access to optimize business. Best of all, eTrust can give you a single view of your entire enterprise, putting you in complete control. As a result, you can make real-time decisions based on comprehensive information. So if you're looking for ways to minimize risk while maximizing your potential, or to get a white paper, go to **ca.com/security**.

ca Computer Associates®

Hello Customers

Radar net flattens Earth for weathermen

■ BY ANN BEDNARZ

Web services, Linux and grid computing are among the technologies researchers are using to develop a system of predicting and improving warning times for weather emergencies such as tornadoes and flash floods.

Last week, University of Massachusetts, Amherst, launched a $40 million research center that will build a nationwide network of radar dishes to collect atmospheric data. The radars would be linked by a grid-computing infrastructure that would let users pool and allocate geographically dispersed system resources as needed.

The Engineering Research Center for Collaborative Adaptive Sensing of the Atmosphere (CASA) says it hopes to overcome a shortcoming of existing weather forecasting and warning systems, which have difficulty monitoring conditions close to the ground because of the curvature of the Earth.

CASA plans to get around the curvature issue and obstructions such as mountains by setting up dense networks of short-range radars that are physically smaller than most existing meteorological radars, says UMass, which is a leader in the CASA project. The radars can be mounted on top of buildings or cell phone towers and supported by PC-sized computers — as opposed to today's high-power radars that often have 30-foot antennas and supercomputer accompaniments.

Over the next few years, a host of municipalities and

CASA is modifying existing radars, such as those used by storm chasers (above), to be networked and installed on cell towers and rooftops.

agencies will install these radar systems. The first field test of CASA will take place in mid-2005 in a tornado-heavy area of Oklahoma, about 20% of the state. Ultimately, CASA plans to place sensors in every city across the country.

"There's computing everywhere," says Jim Kurose, a professor in the computer science department at UMass, of the project architecture. Individual radar sites have Linux-based gear for local storage and computational work, and

larger back-end systems aggregate data from radar sites to run sophisticated meteorological analysis software, Kurose says. The network backbone uses existing infrastructure including Oklahoma's OneNet, a statewide network that provides data and video services to public sector entities such as government agencies, libraries, hospitals and schools.

The demand for real-time forecasting adds complexity to CASA's project. Researchers are designing the system to continually ingest data fed from sensors and use that data to steer the radars toward critical weather conditions, Kurose says. Determining which conditions — such as a hurricane in Puerto Rico or a flood in Houston — receive priority raises complex policy issues.

With lots of end users interested in weather data — including government agencies, emergency response teams and commercial businesses — there might be competing demands for CASA resources, says Daniel Bonelli, vice president of marketing in IBM's Software Group. The research team is devising a system for automating the allocation of computing and storage resources on the fly, Bonelli says.

As a technology partner to CASA, IBM is providing some of the computing infrastructure for the project, including blade servers running Linux, WebSphere application server and integration software, DB2 database and content management software, and Rational application development tools. ■

EDial using IM as hub to integrate voice, data

■ BY JOHN FONTANA

Real-time collaboration is all about integrating various communication tools, and vendor eDial this week will release a server that ties together voice and data using instant messaging as its hub.

EDial's Instant Collaboration

System (ICS) is a standards-based gateway that integrates Web-based instant messaging and presence — a messaging technology that lets users or devices quickly find each other — with telephony, Web conferencing and Web-based document sharing. ICS incorporates Session Initiation Protocol and SIP for Instant Messaging and Presence Leveraging Extensions (SIMPLE), both of which are real-time communication standards.
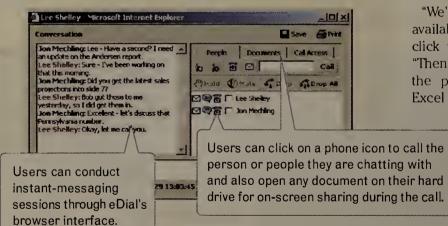
With ICS, eDial beats vendors such as Microsoft and IBM in offering standards-based integration of multi-vendor, real-time communications tools.

ICS uses a browser interface in its built-in instant-messaging and conferencing services, and employs SIP and SIMPLE to integrate other tools, including PBX or other instant-messaging infrastructures on an intranet or extranet.

Within the browser interface, users can create buddy lists, start a Secure Sockets Layers secured instant-messaging session, see who's online or on the phone, click to initiate a single- or multi-party phone call, and begin a document-sharing session or

Real-time collaboration

EDial Instant Collaboration System integrates Web-based instant messaging and presence with telephony, Web conferencing and Web-based document sharing.

Users can conduct instant-messaging sessions through eDial's browser interface.

Users can click on a phone icon to call the person or people they are chatting with and also open any document on their hard drive for on-screen sharing during the call.

Web conference. Call-control features include mute and hold, and sessions are logged and audited.

Reuters is in the final stages of testing ICS throughout its financial news network. It says the browser is key to standardizing the front end of its instant-messaging platform, which it is integrating with instant-messaging services from AOL, IBM/Lotus and Microsoft.

"We are building an infrastructure for collaboration on instant

messaging, and a requirement is to offer the service through a Web interface," says David Gurle, executive vice president for collaboration services at Reuters.

The interface is one of three uses that Reuters has planned for ICS, which will be rolled out to 1,600 users. "We also want to associate presence with telephony," Gurle says. "When our customers want to reach someone they want to know if they are there and if they are available by phone."

Gurle says Reuters also will explore using ICS as a sort of application server for real-time communication to create a link between asynchronous and synchronous collaboration tools.

"We'll be able to see who is available on our buddy lists and click to call them," Gurle says. "Then, instead of explaining over the phone something on an Excel spreadsheet, we can click on our buddy list to start a document-sharing session. ICS is a very elegant way to do this."

Experts say ICS avoids lock-in to one collaboration platform.

"EDial understands the value in tying together these various systems," says David Marshak, director of consulting services for Patricia Seybold Group.

ICS costs $1,000 for 1,000 users. It runs on Windows and Linux, and works with existing voice-over-IP infrastructures. A second version that includes a server for connecting to a legacy PBX costs $8,000 for 1,000 seats. Both options also require a $30 user license per seat. ■

Blades

we needed to add 30 servers. And three, we were out of air conditioning capacity, or very close to it, and we needed to add 30 servers," says Eric French, network manager at the healthcare organization.

"We needed technology that was low-output HVAC [heating, ventilation and air conditioning], had low power requirements and was very small," he says. "Blades fit that bill."

The medical center, which deployed 30 HP BL20p blades earlier this year, is one of a growing number of companies turning to blade servers to get the computing power they need in smaller packages. While blades failed to take off as predicted after blade server pioneers RLX Technologies and Egenera introduced them in 2001, they are gaining momentum, analysts say.

Blade evolution

IDC reports that U.S. blade server sales in the first quarter this year totaled $47 million, eclipsing about $43 million in revenue logged for all of 2002. IDC expects the market is expected to reach $6 billion by 2007.

Blades initially were targeted at service providers and large corporations looking to pack a lot of computing power into small spaces. Today, all the major systems vendors are peddling these slices of processing power as a cost-effective way to consolidate data center infrastructure, get rid of masses of cabling and streamline management. Blades — which are about one-eighth the size of a standard 1U server, but require less power — sit in specialized chassis that enable them to share resources.

Buying one or two blades, however, won't save you money. Because individual blades today are priced about the same as comparable 1U servers and users also must pay

The blade factor

Wondering whether to bring blade servers into your data center? Here are some factors to consider about why or why not to do it.

Why:	Why not:
• **Make room.** An obvious benefit of the diminutive servers is space savings; if you've got expensive real estate, blades might be the way to go.	• **Pared-down processing.** Processing-intensive applications might not be appropriate for blades that typically max out at two processors, though vendors are increasing CPU power.
• **Repeatability.** If you're not running the same configuration on lots of servers, blades will let you get things up and running quickly.	• **Staying attached.** Blades don't yet support all of the same network and switch capabilities as standard servers, though vendors are adding features.
• **Adaptability.** Got spiky loads? Throw in another blade to handle the jump in demand.	• **The numbers game.** Blades make the most sense when you need a bunch; if you need fewer than six, the costs usually outweigh the benefits.
• **Cleaning up.** Dozens of blades mean dozens of CPUs, but without the tangle of cabling.	• **Lack of standards.** The market is still young and so standards have not yet emerged, particularly in the area of management — which means adding proprietary tools to handle blades.

for the blade chassis, which at IBM, for example, starts at $2,800, the savings come only when customers bring in multiple blades, users say.

First-generation blades were simply stripped-down versions of standard servers and had little in the way of additional features, but that is changing as vendors add more intelligent switching, enhanced network connectivity and storage links.

While businesses are looking more seriously at this new breed of server, challenges remain. Early adopters point to benefits such as cost efficiencies, space savings and manageability, but also note that there still is work to be done in areas such as network connectivity, switching and storage.

"There are limitations now, in September 2003, but they will get relieved over time," says Daniel Kaberon, director of computer resource management at Hewitt Associates, a human resources consulting and outsourcing firm in Lincolnshire, Ill.

Hewitt uses grid software on IBM's BladeCenter to spread the load across blades running a pension-benefit calculator engine that provides information on its Web site. Because traffic to the site can spike unexpectedly, he says blades make it easy to meet demands. "As the workload grows, we can simply add more blades," he says.

As for network limitations, "I could talk about limitations in network switches, but I know there are new network switches under development," Kaberon says. HP, IBM and RLX all recently announced network enhancements to their blade servers to make it easier for corporate users to integrate blades into their infrastructures.

It's these kinds of enhancements that have companies considering blades as a more integral part of their data center architectures. Greater Baltimore Medical Center had initial concerns about its blades because the first generation did not connect to storage-area networks (SAN). HP rectified that when it announced connectivity to SANs earlier this year. Today, Greater Baltimore's French says the plan is to standardize on blades.

"Unless there is an absolutely compelling reason not to, meaning there's some application that needs a PCI slot, then it will be a blade," he says.

French says learning to run the blades was about as easy as learning to run any new server.

"The only learning curve, and I think it is probably pretty standard, is getting accustomed to the remote deployment tool. It's one of those things where it takes a little bit of work to get it set up. But once you've got it set up it runs like a champ," he says.

French says he's seeing cost savings, too, especially related to manpower.

"What it used to take two days to do … with the remote-deployment tool and the blades, we can do in two hours," he says. "Every [blade] server has a significant savings in manpower. We're getting to the point where we're ready to turn over [blade management] to our operations staff, and our network people won't even deal with that. [This can] free them up for other projects."

Financial trading systems company Nyfix says it brought in blades at the end of 2001 because it wanted to get a jump on

any technology hurdles that came with the new servers.

"We knew we wanted to adopt blade technology, and we knew there was a steep learning curve, so we wanted to get some experience in operating and configuring blades," says John Knuff, vice president of network engineering at the company in Stamford, Conn. "We're glad we were early adopters because now we've learned some of the tricks of how to use them and when not to use them."

The best place to incorporate blades is when the same configuration is needed across multiple servers, Knuff says.

"Sometimes we bring up several clients a month, and we can scale very quickly," he says. "We can throw in a couple more blades and get them configured and have a new client up and running in a day. The advantage is speed. A disadvantages would be if you want to talk to four different networks on a single server, then I wouldn't use blades."

For Cambridge Health Alliance in Massachusetts, blades meant the ability to quickly add support for a critical ambulatory-care application and to do so on Linux. The alliance's choice represents part of a larger trend. IDC says users deploying blades are doing it at a higher rate on Linux, which represented 15% of the entire server market in the first quarter of 2003 and 57% of the blade server market over the same period.

After determining that buying blades would save the organization $1 million in infrastructure costs over five years, the alliance settled on BladeFrame from Egenera, which hooks into the organization's SAN and automatically moves application load among the blades within the system.

"Our plan is to put smaller applications into logical groups [on BladeFrame] to more efficiently manage the servers and the applications," says Judy Klickstein, CIO and vice president of IT. She says that currently the alliance has small applications on isolated servers that typically run at only 50% capacity. Moving those applications to the blades — where capacity can be shared — would enable the alliance to get more out of its resources, she says. ■

Get more information online.
DocFinder: 7938
www.nwfusion.com

We went with Terabeam's
wireless broadband system
because we couldn't get fiber.
What blew us away was that it delivered the
bandwidth of fiber
at a fraction of the cost.

Terabeam's high-capacity broadband systems deliver the speed, capacity, security and reliability of fiber.

A Terabeam system costs a fraction of fiber, and can be operating in days. Unlike most wireless systems, Terabeam operates at full line rates of up to 1.25 Gbps and is ideal for interference-free, dense deployment. For a free site evaluation or more information call **888-372-2326**.

TERABEAM®

©2003 Terabeam Corporation. Terabeam is a registered trademark of Terabeam Corporation.

Chaos.
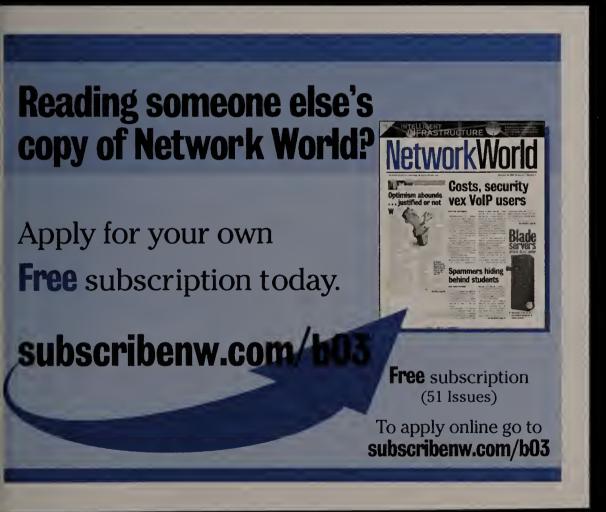
Control.

Take control of your Internet security.

Introducing Proventia™ Enterprise Protection Products. Just because Internet threats are complex, doesn't mean your security has to be. Finally, a single, unified protection appliance that protects more with less, eliminating the cost and chaos of multiple stand-alone security products. Proventia™ centrally-managed products range from detection up to completely unified and proactive multi-function protection appliances, combining firewall, intrusion prevention and anti-virus technologies. Take control of your enterprise security. Switch to Internet Security Systems today. 800-776-2362. www.iss.net/takecontrol.

INTERNET SECURITY SYSTEMS™

Infrastructure

HP to launch high-end net barrage

Desktop Gigabit, 10 Gigabit and 802.11g WLAN gear on tap.

■ BY PHIL HOCHMUTH

HP this week will launch a barrage of high-speed wired and wireless network gear for corporate customers, aimed at

Short Takes

■ **Candera** last week announced a clustered storage controller that joins multiple heterogeneous storage-area networks into a single SAN with a common management interface. The **SCE 510 Cluster** includes two hardware nodes configured in an active-active configuration to protect from failure. With 16 Fibre Channel ports per cluster, they connect to SAN devices and host. The Candera Storage Manager software, which is used to manage the cluster, is a Java-based GUI that runs on Microsoft, Solaris or Linux workstations. With the Candera Storage Manager, administrators can provision, adjust and migrate storage based on user-defined policies. The Candera SCE 510 Cluster is priced starting at $120,000.

■ **IBM** last week introduced services that provide remote access to server computing resources. Aimed at companies that don't want to incur the expense of buying, managing and maintaining their own servers, **Virtual Server Services** lets companies buy server computing power from Big Blue. Customers then pay for what they use. IBM owns and manages the services and keeps them at its data centers. Computing power is delivered remotely to clients. After charging customers a one-time setup fee, IBM bills them according to usage every month. Customers can buy computing capacity on IBM's eServer xSeries machines based on Windows operating systems; eServer iSeries running OS/400; and eServer pSeries AIX-based Unix servers.

delivering Gigabit and secure Wi-Fi links to desktops and 10 Gigabit links in the core.

On tap from HP are new stackable and modular ProCurve LAN switch products that could help companies transition from 10/100M to 10/100/1000M bit/sec desktop connections. For fast backbone links, HP also is releasing its second 10G Ethernet product — a dual-port blade with swappable optics for the 9300m series switch. Also, a new ProCurve 420 Wi-Fi access point supports both 11M and 54M bit/sec wireless speeds, among other features.

HP's new stackable line is the ProCurve Switch 2800 series. The ProCurve Switch 2824 and 2848 include 24 and 48 ports of 10/100/1000M bit/sec Ethernet, respectively. Both switches also include four fiber- and copper-based Gigabit Ethernet ports, supporting up to four uplinks of all copper, fiber or mix of the two. Each port on the switch can auto-sense its connection

speed, from 10M bit/sec Gigabit Ethernet. The switch also supports basic IP routing, which lets the box route traffic among virtual LAN or subnet ports on the same switch without sending traffic to a backbone Layer 3 device for routing, HP says.

Other new Gigabit Ethernet products include a 20-port 10/100/1000M bit/sec module for HP's ProCurve 4100gl series of aggregation switches and a 16-port triple-speed blade for the 5300xl modular switch. Both blades support full IP routing (Layer 3 switching and standard routing protocols) and quality of service with Layer 3/Layer 4-based packet classification and prioritization.

The new 10 Gigabit products for HP's ProCurve 9300m series chassis switch take a different approach than the introductory 10G product the vendor introduced a year ago. Instead of a single-port, fixed-optic

HP's big push

HP is launching new Wi-Fi, Gigabit and 10 Gigabit gear this week. Among the new products on tap are:

● The ProCurve 2800 series 10/100/1000M bit/sec stackable switches for high-speed desktop connections.

● The ProCurve 420, an 802.11g wireless access point, with built-in security and power-over-Ethernet capabilities.

● A dual-port 10G Ethernet module for the ProCurve 9300m series. The new blade offers a choice of optics for short- to long-reach 10G links.

AppIQ boosts storage apps mgmt.

■ BY DENI CONNOR
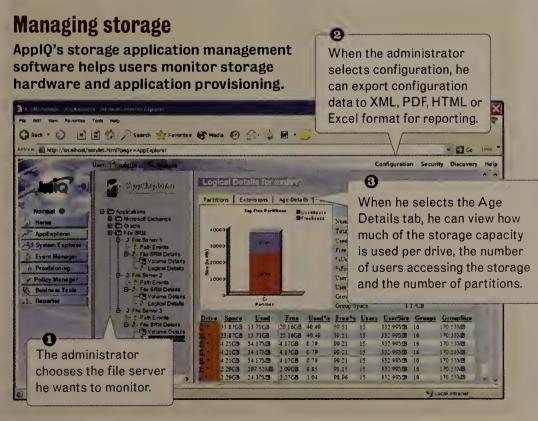
AppIQ last week announced a new version of its storage application management

Managing storage

AppIQ's storage application management software helps users monitor storage hardware and application provisioning.

❶ The administrator chooses the file server he wants to monitor.

❷ When the administrator selects configuration, he can export configuration data to XML, PDF, HTML or Excel format for reporting.

❸ When he selects the Age Details tab, he can view how much of the storage capacity is used per drive, the number of users accessing the storage and the number of partitions.

software that makes it easier for administrators to monitor, manage and report on server and storage hardware and applications.

Renamed StorageAuthority Suite 3.0, the

former AppIQ Solution Suite has been enhanced to include the ability to monitor, report on and manage file servers and instances of Veritas Software's NetBackup back-up and recovery software running in storage environments. Previously, StorageAuthority had modules that monitored Microsoft Exchange and Oracle environments, and provisioned storage required by applications and charged departments or divisions.

Additionally, AppIQ announced that its software now will support Solaris and AIX operating environments, EMC's Symmetrix DMC and Clariion arrays, EMC's Connectrix Fibre Channel switches, and Hitachi Data Systems' multi-pathing Dynamic Link Manager.

A storage administrator would use the StorageAuthority Suite software to adjust or provision the amount of storage allocated to servers running applications such as Oracle or Microsoft Exchange.

Ray Bourgion, vice president of information resources at the Boston Stock Exchange, has installed AppIQ's software to monitor his EMC Symmetrix and direct-attached Sun storage.

The art of the cliché

The Computer & Communications Industry Association recently spent a fair amount of time and money encouraging seven high-profile security gurus to create a 25-page report that boils down to "don't put all your eggs in one basket."

I first heard that sentiment from my grandmother when I was but a wee lad, and it's one I've tried to practice through my personal and business life. But the cliché wasn't enough for the CCIA. The group needed to drag in crop science in the form of "monoculture," theory that holds that it's best to rotate crops in a field. Maybe we should rotate the operating systems in our servers every three years also. Then it co-opted "complexity" theory from applied math, but attempt to pass it off as stating that the more complex a system is, the less secure it can be. If you had to secure a door, would you choose a complex lock or a simple piece of string? Yes, I've hopelessly confused the theory. But, then, so does the CCIA report.

Still, that would only make the report a dozen pages or so, and the CCIA wouldn't publish it without the required 50% content disparaging Microsoft. The more purple the prose, the better. It starts in paragraph two of the very first page: "Microsoft's efforts to design its software in evermore complex ways so as to illegally shut out efforts by others to interoperate or compete with their products has succeeded."

It's helpful to know, by the way, that the major money source for CCIA is Sun.

Still, the sentiment is worthy. Tying all of your productivity to a single thread can be dangerous. Don't put all your eggs in one basket. But don't throw the basket away and certainly don't encourage the government to mandate multiple baskets. And I shouldn't have to say this, especially to seven supposedly dispassionate experts, but it's bad policy to let emotion sway your choices of business tools.

Kearns, a former network administrator, is a freelance writer and consultant in Silicon Valley. He can be reached at wired@vquill.com.

HP

module, the new blade includes two XENPAK-compatible ports, which lets the blade be deployed with different 10G Ethernet optic inserts, such as ports with a ranges of 6 to 24 miles over single-mode fiber. HP says it also will have a multi-mode fiber XENPAK optical port next year for short-range 10 Gigabit links (up to 300 feet). The new two-port 10G blade also will be priced at around half the cost of its previous single-port 10G offering, HP says.

On the Wi-Fi front, HP bills its ProCurve 420 as a "heavy" access point, with security and management features built in. The device is 802.11g-compliant, which lets it communicate with 802.11a (11M bit/sec) or 802.11b (54M bit/sec) Wi-Fi devices. The box also supports the 802.1X end-user authentication protocol, which can lock out untrusted Wi-Fi users at the access point. The device can be powered over an Ethernet cable with support for the 802.3af power-over-Ethernet (POE) standard.

HP announced POE switches in June, but is not shipping them until year-end.

The new access point takes a different approach to Wi-Fi security from HP's previously announced wireless LAN (WLAN) switch strategy — the ProCurve 720 Access Controller and 740 Access Control server, announced in June. Those boxes are designed to centrally control security and management for "light" Wi-Fi access points — inexpensive devices that function as basic 802.11 radios, with network intelligence coming from the WLAN switch.

HP 10G Ethernet products are running in the LAN core at Manchester Community College in Manchester, Conn. Two HP ProCurve 9300m chassis are linked with single-port 10 Gigabit blades in the core, with wiring closet switches at the edge connecting directly to the core boxes.

This two-tier approach is easier to manage than deploying switches in the core, distribution and edge, says Jason Blosser, director of IT for the college.

"This also means all servers connect directly to the core," Blosser adds. "We chose 10 Gigabit because we didn't' want to worry about the core in terms of bandwidth."

The new 10G products from HP could have a place in the school's network in future server consolidation projects. "It would be nice to collapse all of our servers into one large box, then link that to the core with [10 Gigabit]," he says.

Steps such as the boosting of its wired and wireless enterprise gear, and the recent stepping down of HP CEO Carly Fiorina from the board of Cisco, could be signs that HP is preparing to seriously challenge Cisco for networking business in high-end data centers and large corporate LANs.

While HP's high-end server group still lists Cisco as its best-practices partner for networking, that doesn't mean HP won't challenge Cisco with its Gigabit and 10 Gigabit gear, says Zeus Kerravala, an analyst with The Yankee Group. "They may not be a challenge to Cisco right away, but they could become a strong No. 2 in the data center," he says.

The ProCurve Switch 2824 and 2848 are priced at $2,500 and $4,900, respectively. The two-port 10G Ethernet is priced at $35,700, with single-mode-fiber Xenpak optical inserts costing $13,000 each. The 16-port 10/100/1000 blade for the 4300gl will cost $2,200. The 420 Wi-Fi access point costs $470. All of the new HP products will ship this quarter. ■

ApplQ

"[The software] allows us to have a single interface into the multiple different storage devices we have within the facility here," Bourgion says. "As we move forward, we will tailor the storage environment to the specific needs of the application so that potentially business-critical applications and data will sit on the EMC Symmetrix, and the less-critical applications will be put on storage that gives us a better price point."

StorageAuthority for File Servers is a module that lets the suite report on space utilization; the type, age or size of a file; users who are exceeding their disk quota; and the arrays to which a file server is assigned. In addition, StorageAuthority for File Servers can manage the provisioning of storage for a file server. It works with Solaris, Windows, NetWare and AIX file servers.

The StorageAuthority for Veritas Net-Backup module includes the ability to monitor the success of full or incremental backups, discover and visualize back-up resources, and warn of and report errors.

While ApplQ is not alone in the storage resource management market, its product is differentiated from others such as the Creekpath Suite or Veritas' SANpoint Control by its ability to link storage allocation to application requirements.

"ApplQ has a great example of a storage management console — it's an emerging part of the market that administrators can use to collect and monitor information across your storage environment from the host to the back-end disk," says Jamie Gruener, an analyst with The Yankee Group.

"ApplQ organizes the environment by tying the application data to the storage system and manages it in a way that the application is fully involved in the management," he says.

In addition, StorageAuthority is among the first software that fully supports the new Storage Management Initiative Specification (SMI-S). SMI-S is an industry-wide project and specification to enable storage software and hardware with an object-oriented storage management framework, the Common Information Model and Web-Based Enterprise Management, which establishes the ability to share CIM data between devices and software. ■
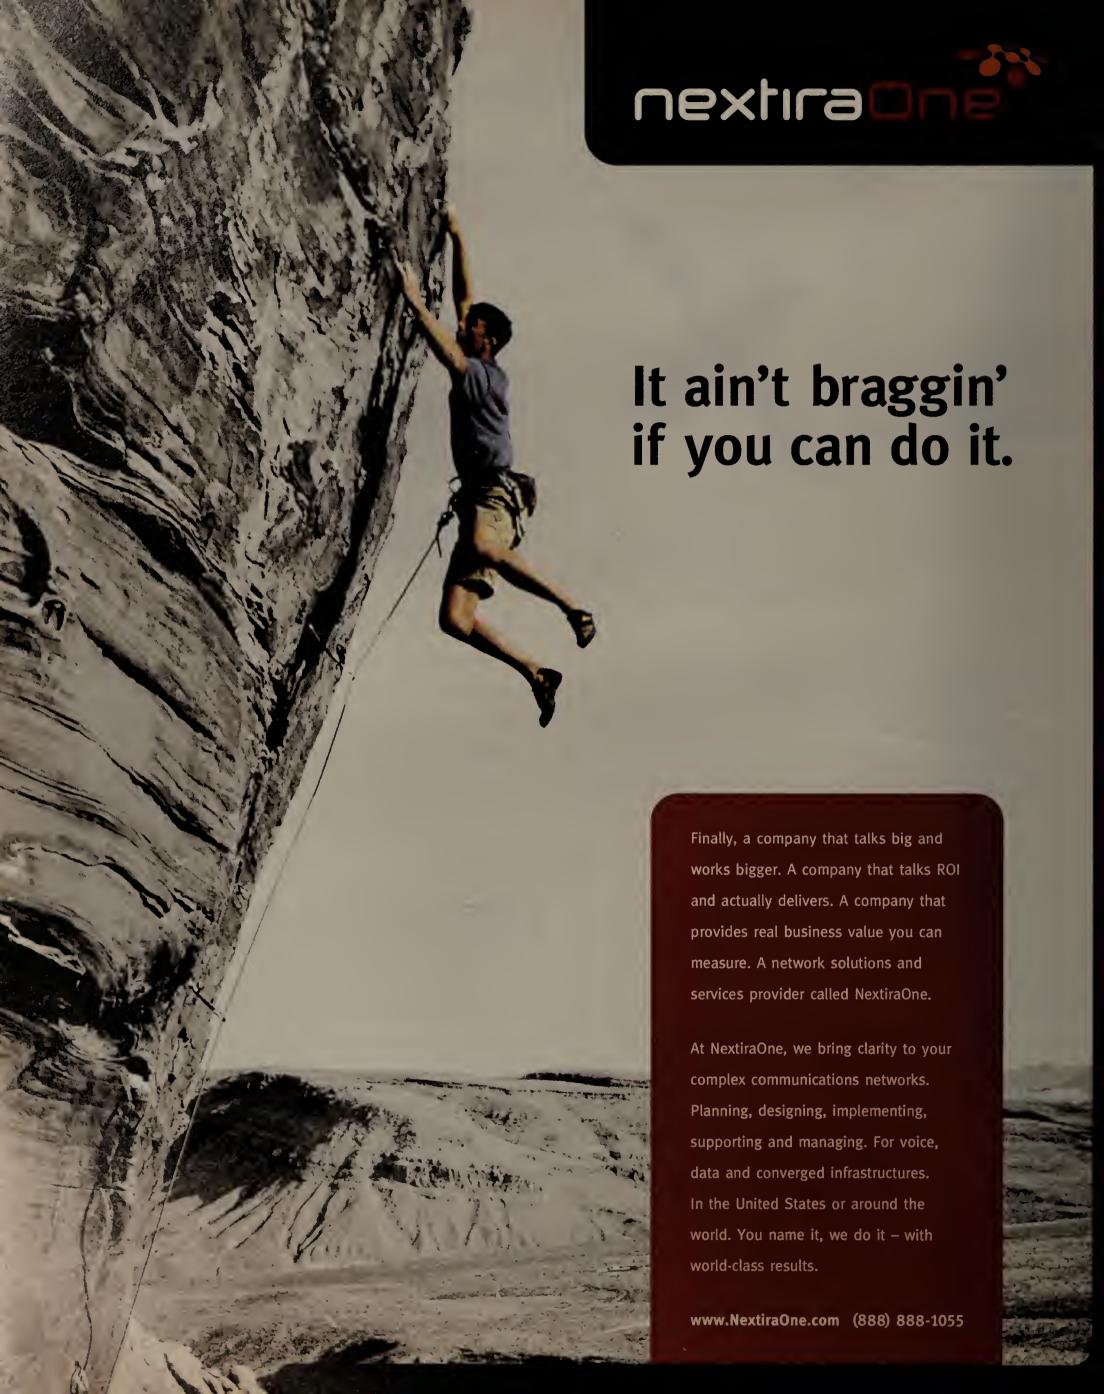
Scrap the multiple business telecom providers. Scrap the confusing multiple rates. In fact, scrap telecom as you know it. With XOptions,™ you get everything your business needs for one flat monthly rate.* On one simple bill. No unhappy surprises. No mind-boggling rate structures. Just a total communication package of local, long distance, Internet and Web hosting. So long hassles, hello simplicity.

XO
Not just talk.

nextiraOne

It ain't braggin' if you can do it.

Finally, a company that talks big and works bigger. A company that talks ROI and actually delivers. A company that provides real business value you can measure. A network solutions and services provider called NextiraOne.

At NextiraOne, we bring clarity to your complex communications networks. Planning, designing, implementing, supporting and managing. For voice, data and converged infrastructures. In the United States or around the world. You name it, we do it – with world-class results.

www.NextiraOne.com   (888) 888-1055

Enterprise Applications

ID mgmt. fuels roles, rules growth

■ BY JOHN FONTANA

In these days of distributed networks, user management is not for the faint of heart, and that is increasing interest in two techniques for streamlining the process.

Roles and rules are two approaches that promise automation and efficiencies in provisioning resources to users, and consistency in granting and revoking access rights. The goal is to replace the error-prone manual process of performing those tasks one user at a time with what amounts to batch processing.

Using roles- and rules-based models can help tighten security of network resources and ensure compliance with federal regulations such as the Sarbanes-Oxley Act, Gramm-Leach-Bliley Act and Health Insurance Portability and Accountability Act.

Roles are predetermined sets of access privileges that are associated with a group of users on a network. Users are assigned to roles. The National Institute of Standards and Technology (NIST) developed the model, called Roles Based Access Control (RBAC), more than a decade ago. The Massachusetts Institute of Technology, Stanford University, Sun and Pricewaterhouse-Coopers are among those that have developed their own roles-based models.

In comparison, rules were introduced recently with the advent of provisioning systems. They are more flexible and act as "if/then" expressions that are executed within software when a user attempts to access a network resource. For example, a rule might state "if" the user has the title "sales manager" and works in Division A "then" he is entitled to access System B.

Combo pack

Experts say that a combination of the two might be the best approach in meeting today's requirements for identity management.

"We found that just using roles would not be enough to provision users," says Steve Linstead, directory services architect for Johnson Controls, a Milwaukee supplier of automotive parts and building controls, such as heating/cooling.

Johnson Controls is finishing a pilot project with provisioning software from Netegrity that will be implemented next year. "Roles left too many gaps, and we needed rules to further define the user. We can have a supervisor role, but supervisor of what? The rule then determines how the role operates," Linstead says.

Interest in roles and rules is accelerating, especially with the number of networked applications growing along with the internal and external users seeking access. Corporate users are seeking options, and vendors such as Beta Systems, Business Layers, IBM, Microsoft, Netegrity, Novell, OpenNetwork Technologies, RSA Security, Siemens and Waveset Technologies are listening.

"Most companies today are under pressure to do more with roles- and rules-based user management," says Christy Hudgins, president of Hudgins Group, a research firm. "I see differing motivators among different types of businesses. Some retailers are very cost-reduction-driven, while others are most interested in relieving the administrative load on IT staff. Regulatory compliance is a big factor with regulated financial institutions, as well as medical groups. Security tends to be the big driver with retail banks."

Hurdles to clear

However, the road to exploiting efficiencies using roles and rules is paved with scalability problems and complexity in defining roles and rules that align with business processes, such as creating new user accounts.

Experts say users must be cautious when implementing roles and rules, which is most often done through provisioning, access management or directory software.

They say XML-based policy languages eventually will further combine with roles and rules for user management among corporate networks integrated through Web services.

"If you have thousands of people needing access to your network because of a hiring cycle or contract work, all those accounts are set up with the right level of access, authorization is done once and done consistently, and there is less opportunity for human error," says David Shapiro, assistant director, Americas IT for Ernst & Young. The company has used roles and rules within its provisioning software for the past six years to set up new users with necessities such as network access, telephones, building security badges and business cards. "We don't have people going to each server to set up accounts. What we have is a repeatable business flow, a workflow to support that process."

Experts say defining those processes is key.

"Coming up with role definition is hard work," says Gerry Gebel, an analyst with Burton Group. "Rule definition also has a similar process." But best practices such as limiting the initial scope of the project and getting people involved from business managers to IT helps immensely.

Once roles are deployed, the work is just beginning, Gebel says. Auditing must be

Ease of use

The concepts of roles and rules can help corporations establish effective provisioning and access controls for large groups of users, but upfront planning and an understanding of how the two concepts can complement each other and how they contrast can be the difference between success and disaster.

	Pros	Cons
Roles	• Suitable for many scenarios. • Conceptually simple. • Proven in large scale ERP implementations. • Mandated by some regulations. • RBAC model well known.	• Project scope must be defined carefully. • Defining roles can be an expensive and lengthy process. • Difficult to engineer a model across an organization that is dynamic and has few static, well-defined roles.
Rules	• Intuitive and clear interpretation. • Easy to write for simple domains.	• Limited without using roles. • Challenging to update and keep consistent. • Change management a necessity. • No standards.

Short Takes

■ Neoteris last week improved its software access-management controls and released a host of new features to secure content accessed via a browser. With Version 3.3 of its **Instant Virtual Extranet** platform, Neoteris has added single sign-on capabilities, integration with password management software and tools to check processes running on remote devices. IVE 3.3 software will be included in the Neoteris Access Series and Meeting Series products. The new **Host Checker Agent 2.0** looks for software compliance and monitors the executable processes running on a target machine to ensure it is not malicious. The agent can shut down a session if it detects something out of the ordinary. Neoteris also has added a **Cache Cleaner Agent,** which cleans out temporary files and cookies stored on a browser. Also new is **In-Transit Data Protection,** which allows non-cacheable Hypertext Markup Language rendering to ensure data is not left behind on the client software. The company has added native single sign-on technology to IVE, including form- and header-based exchange of user name, credentials and other attributes. The Access Series, which includes IVE 3.3, starts at $10,000. The Meeting Series upgrade to the software costs $2,000.
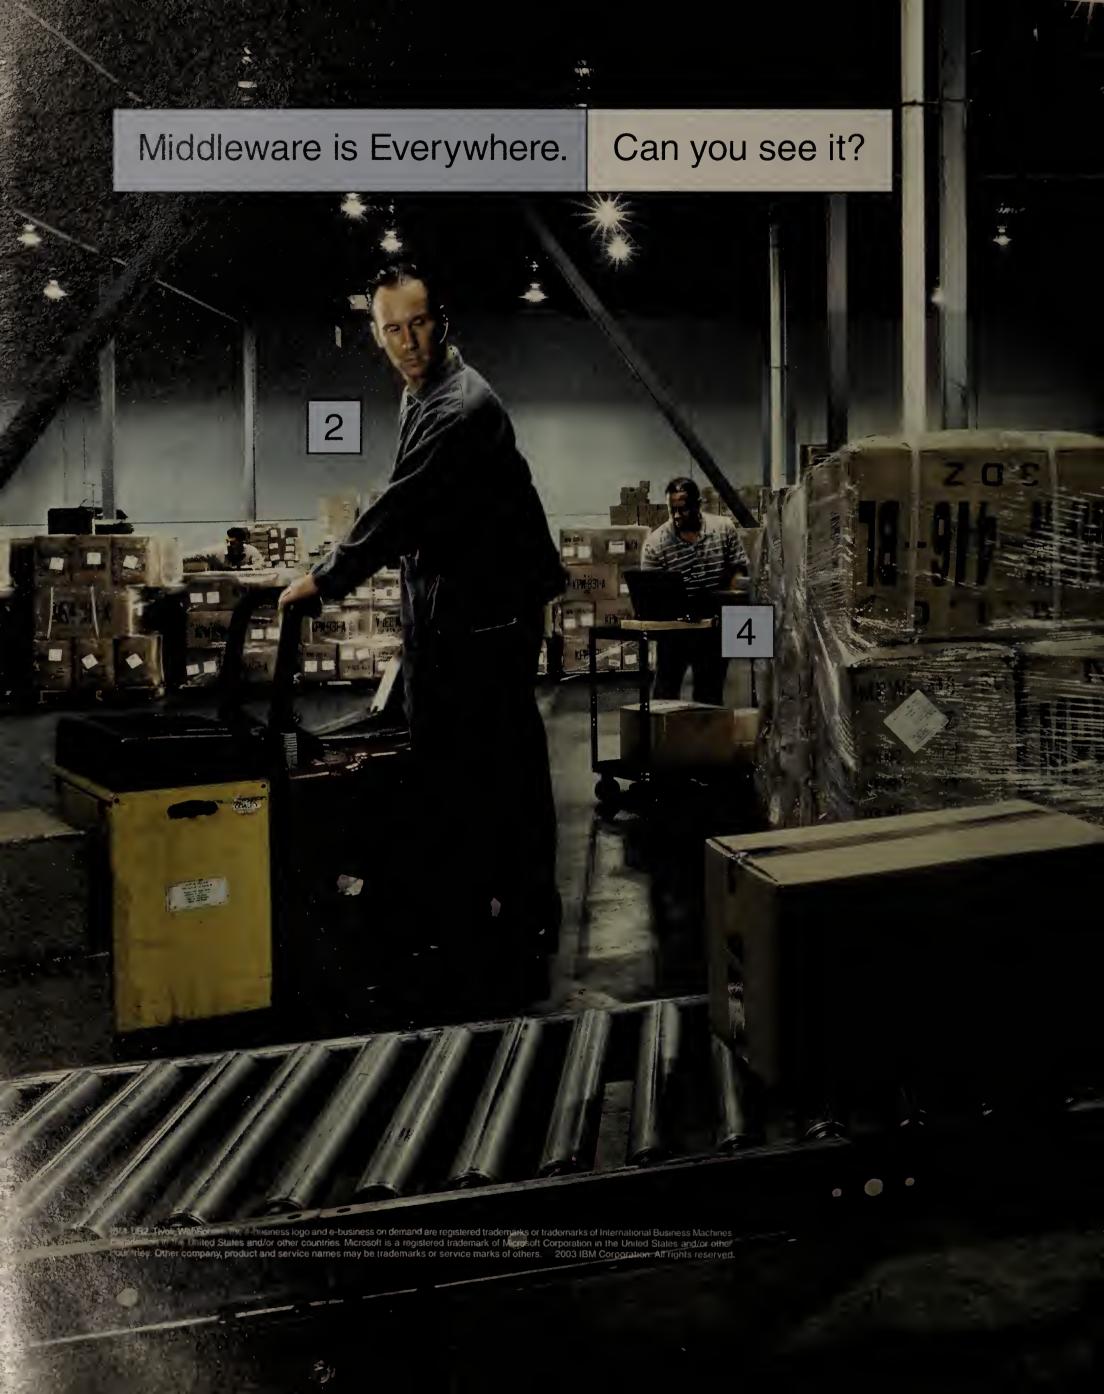
■ Application-security appliance vendor **Teros** has announced Version 3.0 of its **application firewall,** adding protection against cross-site scripting attacks that let hackers steal computer users' desktop cookies. The software features new defenses against denial-of-service attacks and a "cloaking" mechanism to camouflage information about corporate domain naming. Teros 3.0 is available for $25,000.

Middleware is Everywhere. Can you see it?

KEY

1. Automatic overview of operation.
2. Automatic shipping of sale.
3. Automatic identity verification.
4. Automatic updating of inventory.
5. Automatic tracking of delivery.

MIDDLEWARE makes the on demand world on demand. And middleware is powerful IBM software like Tivoli,® DB2® and WebSphere.® Open, behind-the-glass technology that can automate it all – IBM, Microsoft,® Oracle. Problems are foreseen and solved before they occur. IT resources are directed to core business needs. Costs are reduced. It's automation. On demand. And it makes your customers happy. Very happy. @business on demand.™ Go to **ibm.com**/software/automate

Blow-by-blow coverage

I'm a motor racing fan. Well, at least a fan of some types of motor racing. Dirt track, Figure 8 and the IRL do not do that much for me. But Formula 1, CART, Le Mans style endurance, Isle of Man TT and NASCAR racing get my attention during the season, and I get a touch of withdrawal during the winter when most racing is gone from the tube. I have found that the experience of being a motor racing fan has changed dramatically over the last few years as motor racing has embraced the Internet.

Most major racing series now can be followed in real time on the 'Net in a level of detail that rivals what the booth announcers have access to. Formula 1 and the American Le Mans Series have Web sites that provide auto-updating Web pages that show the position of every car, the gap to the car in front, lap speed and other details. NASCAR has the same sort of thing but you have to pay for it, and the service does not support Macs, so I do not use it.

A major result of these services for me is that I generally do not turn on the TV sound. I prefer to listen to KHYI (country music streamed over the Internet) instead. For some reason, most TV commentators have not figured out that TV is a visual medium and insist on describing what I'm already seeing, and they seem to feel that a second without babble is a second wasted (or maybe they are paid by the word). The Internet coverage also continues through the large number of increasingly dumb commercials. (I fail to understand how showing two to four times an hour that Dell hires imbeciles as interns is supposed to make me want to do business with them. For me it does the reverse. Why should I buy something from a company whose only visible employees are too stupid to turn on a light switch and so dishonest that they lie about why the light is off?)

Motor racing is not the only sport that has discovered the Internet. You now can get live scores and stats for seven of the top 10 most-hated sports (www.nwfusion.com, DocFinder: 7926). At least I've not seen Web sites that purport to provide live coverage of dog fighting, pro wrestling or bull fighting — although I suspect that there might be sites covering bull fighting in some countries.

I cannot predict what the general effect will be of this Internetization of sports. But I have a harder time working at the same time as a race is on the Internet than I did when it was just on the tube. I watch the action and the stats more closely on the Web and see the ads on the cars (of which there are many) more than I do with NASCAR, for example. NASCAR has made an error in charging (and using non-standard technology) on its site. The organization would attract more viewers if it enabled access to the basic stats for free and only charged for the fancy extras. Time will tell if the general model is open or closed.

Disclaimer: Harvard, to some, is a fancy extra, but I'm the one following F1, not the university.

Bradner is a consultant with Harvard University's University Information Systems. He can be reached at sob@sobco.com.

Roles and rules

done to keep role and rule definitions up to date, a process that is tougher with rules because they have more data and policy information than roles.

Vendor Eurikefy offers tools to help define roles and rules, and audit them once they are deployed to ensure they are correct.

Refinement is ongoing

NIST is working on improving RBAC, by performing tasks such as creating dynamic roles that include characteristics similar to rules, improving it with workflow capabilities and integrating it with Web services applications. NIST also has submitted RBAC to the American National Standards Institute for adoption.

Rules have their own sets of challenges, including the need for standardization. Web services protocols such as Extensible Access Control Markup Language and the emerging WS-Policy should begin to erase that limitation.

Still there are others who say roles and rules need even more help.

"Rules came about when the limitations with roles hit," says Vivek Pabby, vice president of applications development at The Depository Trust & Clearing Corporation (DTCC), the largest financial services post-trade firm in the world. "You have to track and administer rules, but there is no auditing or security associated with rules, and it becomes a maintenance nightmare." Pabby says roles and rules are two-dimensional and that another layer is needed that puts user management into context.

DTCC uses software called Concero from TruLogica that is billed as context-based user management.

Concero uses three constructs: a service, which defines an application; groups, which are sets of users; and business relationships, which define exactly what parts of an application a user can access.

But it also goes a step further to incorporate approval workflows and policies in the service. Concero also includes the ability to delegate user administration internally and with external partners to allow scalability, an issue that has hampered adoption of RBAC.

Pabby says user registration that used to take up to 10 days now is done in real time or within 24 hours if an approval is needed. Account termination, password reset and audits all happen in real time instead of days or weeks.

"Roles and rules can work for smaller organizations that only operate within their boundaries," Pabby says. "But we have 4,000 external partners, and roles and rules don't meet the requirements for our environment." ∎

Adding business smarts to service desk savvy

■ BY DENISE DUBIE

When John Bandy decided to replace Foremost Farms' homegrown, document-based help desk software with new, more automated features in HP's OpenView Service Desk software, he got a pleasant surprise: a $60,000 savings in network uptime.

For Bandy, the IS technical manager for the dairy cooperative in Baraboo, Wis., the new software would reduce the four days it usually took to resolve one trouble ticket to a little more than a day and a half. Also, and possibly more important to corporate management, the service desk helped Bandy avoid network outages, which would have cost the company between $50,000 and $60,000. He estimates that an hour of downtime costs Foremost Farms $1,200, and with no outages in the first year of using the software, he proved to management the software investment paid for itself.

"We can now prove that we have fewer errors and show the importance of implementing good processes to business managers and prove our worth as a department," Bandy says. "We wanted to re-engineer our processes and add more accountability, follow-up and business processes into our service desk."

Larger trend

Bandy's moves are part of a larger trend toward automating mundane network management tasks such as trouble ticketing. A slew of vendors, including Computer Associates, HP, Peregrine Systems and Remedy (now owned by BMC Software) and Peregrine Systems, is powering the typical trouble-ticketing tools with more automation, business process mapping and integrated management features.

Service desk software, which streamlines the process of tracking service problems and following them through to resolution, isn't just about kicking off trouble tickets and logging calls to the help desk anymore.

"People aren't really talking about tracking tickets. Instead they want to use [service desk software] to operate IT better as an organization," says Jasmine Noel, principal analyst at JNoel Associates. "Putting that grease in the gears between IT departments and business units can result in big operational efficiencies, which turn into big cost savings."

An average help desk could receive a mix of automatically generated trouble tickets (which are triggered by a network or application failure), manually documented problems input by IT staff and a variety of end-user support calls ranging from complaints about a slow application to logon and password resets. But today's help desk offerings try to go further to speed problem resolution and response to end-user calls.

Automation features in the software can detect potential problems earlier, before users are affected. Self-service management portals provide end users with a simple answer to their more basic help desk questions. Consolidating tools such as systems and asset management software with service desk products also can give network managers a more complete view of the network they serve. And new integration capabilities and documentation features make it easier for network managers to incorporate business processes into the IT strategy.

Trends and directions

Service desk software can help lower costs, reduce network and application downtime, and improve overall end-user satisfaction — upping the IT department's credibility. Here are the key components of an effective implementation:

- **Automation:** Incorporating automation can more quickly generate, assign ownership and resolve trouble tickets on enterprise networks.

- **Mapping business processes:** Documenting processes can help service desk systems alert IT and prevent a network or application failure from affecting end users.

- **Self-service portals:** Help desk duties such as password reset present perfect portal opportunities and eliminate end user calls for simple tasks.

  Change management: Monitoring change on network devices can help IT staff more quickly determine the source of problems — for example, if a recent configuration change caused a router to misdirect traffic.

- **Consolidation:** Tying the service desk into systems management tools, business process monitors and network asset monitoring software will give help desk staff a more complete picture of the network.

Manual labor

Garter estimates that, in the worst-case scenario, companies without automated help desk software service would spend

$930

per user, per year; best case would be about

$130.

Service desk software can be more about processes than product. Bandy signed on with HP because the company committed to following the Information Technology Infrastructure Library's standards, a set of best practices for operating and implementing IT in companies. With good processes in place, software tools can automate tasks and improve operational efficiencies. But without the processes, the software won't give results, Bandy says.

Vendor offerings vary, but corporations use the help desk software to define who owns which problems across business units, and then an administrator is responsible to log the actions taken to resolve problems.

Following a distributed model, centralized server software houses the rules and policies established and communicates with managed devices via network protocols or collects data from software agents installed across the network.

Products can track assets, changes and frequency or patterns of recurring problems to help IT staff fix the infrastructure and avoid

future failures that could cause network downtime.

Remedy has addressed such a need with its recent release of IT Service Management Version 5.5, which includes more features in its Asset Management, Service Level Agreement, Help Desk and Change Management products. Plus, Remedy is adding workflow among these applications. The company can take advantage of service modeling technology acquired by parent BMC, which purchased IT Masters earlier this year and incorporated its MasterCell modeling technology across product lines.

"We can provide one field now that says, 'This is the business service being affected,'" says Rick Fitz, director of product marketing and management for IT service management at Remedy. "That visibility helps to eliminate some of the noise when poring through alerts and tickets."

Help for the help desk

The help desk team at Schaller Anderson, a national healthcare management and consulting company in Phoenix, put Remedy service management software into production in April. Terry Newman, director of IT, says the approximately 1,300 users would generate about 60 trouble tickets per day. On average, it took the help desk about eight hours to resolve problems, and Newman reports his staff cut that time in half on about 70% of the tickets. He credits understanding Schaller Anderson's business units and ultimately the company's end users for the software implementation's speedy results.

"We met with managers of critical areas because it's our job to understand enough about the business needs to prioritize them in the software," Newman says. "Every problem is the most important problem to end users."

Ken Hooky, director of IT for the Royal Bank of Canada in Toronto, didn't have much of a choice when his organization first started working with Peregrine and its service desk offerings. The bank, formerly a customer of IBM Tivoli, migrated to Peregrine when Tivoli sold its Service Desk suite to Peregrine. But he says the software now helps his team manage about 2,000 problem tickets and 1,500 call tickets per day. The secret, Hooky says, is to track processes.

"We deal with every business unit differently, and we don't necessarily track the entire infrastructure," Hooky says. "They define their processes to us, and we map them into the software."

While he'd like to see Peregrine make the product more user-friendly with a better front-end user interface, Hooky says the basics are there.

"The key is to provide flexibility in the software because service at different companies will vary," Hooky says. ■

Network Management
Subscribe to our free newsletter.
DocFinder: 5434 www.nwfusion.com

The

"STOP THE
BLAME
GAME"

storage software company.

Service Providers

AT&T expands Ethernet MAN offering

Carrier turns to new strategy of provisioning on demand rather than building everywhere.

■ BY DENISE PAPPALARDO

A new metropolitan Ethernet service that AT&T launched last week lets users go beyond point-to-point LAN connectivity.

Short Takes

■ **Verizon Wireless** last week launched a commercial data service in **Washington, D.C.,** and **San Diego** that usually will deliver several times the speed of a dial-up connection. The services will offer average rates of 300K to 500K bit/sec, with burst rates up to 2M bit/sec. The service will cost $80 per month for unlimited use.

■ **NaviSite,** which added hosted messaging services to its repertoire when it acquired Interliant earlier this year, offers a shared managed messaging service that the company says is designed to provide the reliability of its dedicated service but at a lower price. The fully managed, shared **Microsoft Exchange 2003 hosting offer** includes migration services to help companies upgrade from Exchange 5.5, Exchange 2000 or other platforms, NaviSite executives say. Users pay a monthly fee, which ranges from $12 to $15 depending on customer configuration, and gain all of Exchange 2003's functionality.

■ **ADC** and **Colubris Networks** have signed an OEM agreement whereby ADC will offer Colubris' **Wi-Fi hotspot** products to wireline carriers. Under the agreement, the companies are developing an integrated Wi-Fi/DSL access device for incumbent local exchange carriers. Support for Wi-Fi and DSL in one device would let carriers utilize their copper infrastructure and existing network assets as they extend broadband service to wireless networks, the companies say. Two regional Bell operating companies will begin lab trials of the new product later this year.

The carrier's Ethernet Switched Service Metropolitan Area Network lets customers connect multiple LANs without the cost of provisioning dedicated connections between each site.

Users can chose from 50M bit/sec up to 1G bit/sec dedicated Ethernet connections to AT&T's metropolitan network. The carrier is guaranteeing network availability from 99.9% to 99.99% depending on how it provisions each connection.

AT&T says the service is available in 67 metropolitan areas, but that does not mean network gear is deployed to support the service today. The carrier has moved away from the "if we build it they will come" philosophy of rolling out a new service.

"[AT&T] is not spending the capital upfront to deploy the service" in all 67 metropolitan areas, says Franco Callocchia, director of Ethernet services at AT&T. "We can deploy a network on demand."

In some cases, AT&T will have to light fiber to connect customers to its local network, and that could require up to 90 days. Customers with offices in a building with an existing fiber-optic connection to AT&T's local network could have

services within days.

"It's smart not to deploy gear without knowing how much demand is out there," says Sterling Perrin, an analyst at IDC. The failure of Ethernet service providers in the past could be attributed to the fact that they spent millions to build networks and demand waned, he says.

AT&T is provisioning its Ethernet Switched Service customers in three ways: over existing fiber, over its SONET infrastructure or using one of its multiple Ethernet service provider partners.

Excel Switching cuts the wires

Programmable switch takes on mobile network duty.

■ BY JIM DUFFY

Excel Switching, a maker of programmable call control and media processing switches, is going wireless.

The company is adding support for wireless protocols to its Converged Services Platform (CSP) switch. This extension lets service providers offer applications and services that span second-generation, 3G and Wi-Fi wireless networks that support Signaling Systems 7, Session Initiation Protocol (SIP) and wireless protocol signaling, Excel says.

The CSP can be configured as a proprietary service node or an standards-based Intelligent Network node. The wireless Intelligent Network protocols the CSP now supports include:

● ANSI-41, a messaging protocol used in Code Division Multiple Access and Time Division Multiple Access networks for intersystem hand-off, automatic roaming, authentication and supplementary call features.

● Wireless Intelligent Network, another messaging protocol that lets subscribers to an ANSI-41 mobile network access certain features while roaming.

● Mobile Applications Part, a messaging protocol used in GSM networks for user authentication, equipment identification and roaming.

● Customized Applications for Mobile Networks Enhanced Logic, which adds Intelligent Network functions to GSM networks so a subscriber's "home" network can monitor and control calls while the sub-

Excel's latest

Features of Excel's Converged Services Platform for wireless networks include:

● Execution of service applications that can span 2G, 3G and Wi-Fi networks.

● Adaptable between service node and wireless Intelligent Network node.

● Single API that supports development of multiple wireless service applications.

● Extendable to future wireless applications, protocols and network architectures.

scriber roams.

These protocols let service providers develop new CSP-hosted wireless applications and services such as prepaid calling, mobile Centrex, personal access/follow me, Short Messaging Service, presence, operator services and intelligent call screening. Carriers who use the CSP to develop and deploy wireline services now can use the same switch to deliver wireless applications and interconnect public switched telephone network and IP infrastructures.

Analysts say support for wireless proto-

cols, applications and services is a natural evolution for the CSP, which is installed in 125 customer sites worldwide.

"Wireless is an area where operators are generating revenue-producing services beyond simple transport," says Dave Passmore, research director at Burton Group. "It kind of makes you wonder why Excel didn't do this sooner."

The company will face competition from several start-ups, such as Megisto Systems, Tahoe Networks and Watercove Networks, but these companies are struggling to gain momentum, Passmore says.

"Excel has a much easier job with its pre-existing relationships [with carriers] and a wide range of protocols," he says.
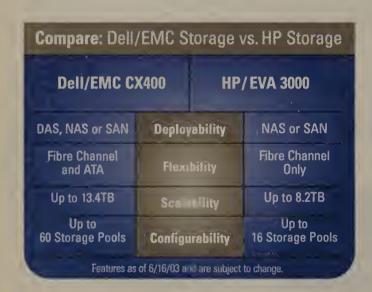
Excel's challenge will be in managing the transition to IP, Passmore says.

"Right now, they're sitting pretty. But in a couple of years, systems based on the old [Intelligent Network protocols] may come tumbling down when everything moves to SIP and IP," he says.

The wireless capabilities for the CSP are available now. Pricing for the CSP also was not available. ■

Dell/EMC SAN solutions.
More room,
less rent.

Dell | EMC Storage Solutions

Large or small, your company can have a flexible storage solution from Dell. As you can see, Dell offers a variety of solutions that give you the flexibility to grow. Complete storage solutions – including software and services – that deliver maximum productivity and scalability.

See for yourself why companies from small business to the Fortune 500 are turning to Dell/EMC SAN solutions. Go to www.dell.com/SAN4 today and click the Storage Consolidation ROI Analyst Tool.

Compare: Dell/EMC Storage vs. HP Storage

Dell/EMC CX400		HP/EVA 3000
DAS, NAS or SAN	Deployability	NAS or SAN
Fibre Channel and ATA	Flexibility	Fibre Channel Only
Up to 13.4TB	Scalability	Up to 8.2TB
Up to 60 Storage Pools	Configurability	Up to 16 Storage Pools

Features as of 6/16/03 and are subject to change.

The Dell/EMC CX600   The Dell/EMC CX400   The Dell/EMC CX200

EMC^2

Complete SAN solutions at a better overall value. Easy as **DELL**™

Click www.dell.com/SAN4 Call 1-866-871-9877
toll free

EYE ON THE CARRIERS
Johna Till Johnson

IT executives think a lot about disaster recovery these days. And that's a good thing. As companies continue to consolidate resources — data centers, facilities, networks — the need for redundancy becomes critical.

All too often, though, in their rush to back up their systems and facilities, executives overlook back-up strategies for their networks.

That's bad news. In today's distributed environment, losing network connectivity can short-circuit a company's ability to function just as effectively as a massive system or data center outage.

Here are some best practices for designing and developing a resilient WAN:

● Ensure physical redundancy. Make sure you have physical diversity in your cable runs (even if the circuits are provided by different carriers, you might need to double-check that the providers aren't sharing strands of the same cable).

● Ensure logical redundancy. Check that

Tips for designing a resilient WAN

your IP services have alternate routes. Multihome your IP links or set up a redundant connection to an IP network access point served by multiple IP providers.

● Check for carrier facility and power redundancy. It doesn't do any good to connect to a carrier POP if the switches are down because the power is out. Make sure your service providers have back-up sources (including diesel generators).

● Protect remote offices and workers. Don't assume everything's fine just because your site-to-site connectivity is in place. Many times, remote offices lose all functionality if they can't connect back to centralized data and applications. Look into alternate technologies to provide this connectivity, including dial-up, home broadband (DSL and cable modem) and, increasingly, Wi-Fi. IP VPNs also can provide a highly effective, low-cost back-up strategy.

What about voice? Make sure your voice network is backed up as effectively as your data WAN.

If cell phones are your back-up mechanism, make sure you have an up-to-date directory of numbers. And does every employee at your company have a cell phone? For some organizations (grocery stores and other retail organizations, for instance) equipping every employee with

a cell phone is impractical. For outgoing calls, calling cards might be an option (assuming your employees can reach a functioning phone). How do you plan to handle incoming calls?

If you've implemented IP telephony, make sure you dedicate the same care to backing up your IP telephony system as you have for your voice and data networks. Be particularly sensitive to power issues.

Test regularly. Not long ago, I asked a group of IT executives how often they performed a full soup-to-nuts test of their back-up plans, including facilities, systems, and networks. The answer? Never. Bad news. Yes, it's difficult. Take the time, and figure out a plan. You won't be sorry.

Finally, develop an effective communications plan. Figure out how to alert folks in the event of an outage, or even a test. If you need them to change their business processes, how will they know to do so?

Bottom line: When you're thinking of ways to make your data centers, servers, and systems more redundant, don't forget the network.
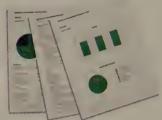
Johnson is president and chief research officer at Nemertes Research, an independent technology research firm. She can be reached at johna@nemertes.com.

AT&T

The carrier also would not elaborate on which Ethernet service providers it's partnering with to support this Ethernet service or any of its other Ethernet offerings that it has rolled out in the past few years (see graphic, right).

This use of multiple provisioning methods, all of which support Ethernet over fiber or SONET, is why the carrier says it cannot provide standard pricing or service-level agreements (SLA) at this point. But Callocchia says AT&T will move to standard SLAs and pricing over time.

AT&T's Ethernet offering is an alternative to standard OC-3 to OC-12 private-line services. The service provider is competing directly with incumbent local exchange carriers, competitive local exchange carriers and niche Ethernet providers.

Familiarity important
The carrier's Ethernet Switched Service offers simplicity because it is based on a protocol more users are familiar with, Callocchia says. Users connect to the service with a standard 10/100M bit/sec Ethernet port, which means customers can use the same router they already have deployed to link up to AT&T's local network.

The carrier also says users should expect cost savings, but that's difficult to substantiate because the carrier did not provide pricing details. But Callocchia says AT&T's Ethernet service is priced competitively compared with other Ethernet offerings on the market.

This is strictly a metropolitan offering, which only lets users connect sites within one region.

The carrier says it will offer additional WAN Ethernet offerings in early 2005. Callocchia says these services will support local Ethernet connectivity for IP VPN, ATM and frame relay customers. AT&T is testing these services. ■

Ethernet from AT&T

AT&T started rolling out Ethernet services in 2001. Here are the carrier's current and future options:

WAN access options
● Managed Internet access service Ethernet access.
● Ethernet Access to IP VPN service (planned for 2005).
● Ethernet Access to ATM/Frame relay service (planned for 2005).

Metro options
● Ethernet private-line service MAN – Point-to-point Local.
● Ethernet private-line service MAN – Point-to-point Long Haul.
● Ethernet Switched Service MAN – Any-to-any LAN connectivity (announced Sept. 27).

SONET alternatives
● Accu-Ring Ethernet service channels.
● Ultravailable Managed OptEring service.

Hughes expands high-speed Internet offer

■ BY DENISE PAPPALARDO

Hughes Network Systems last week announced an enhanced version of its satellite Internet access service.

HNS says its Direcway DW6000 terminal would let users support multiple PCs or Macintosh computers in a home office. Previous versions limited users to connecting a single PC to a terminal. For the first time, users now can connect Macs to the service, but only if they deploy the service provider's new terminal.

Direcway is an alternative to DSL or cable modem broadband services, but only in areas where these are not available. While Direcway offers a similar service that lets users jump from dial-up to a broadband service, it's more expensive than both.

The standard service costs $60 per month, which allows for two simultaneous broadband connections. Business users who want a static IP address and dial-up access can pay $80 per month.

While the monthly service charges are slightly more expensive than DSL or cable, the equipment costs are a lot more. HNS charges $600 for its DW6000 terminal. Many DSL providers offer customer premises equipment at no cost.

The service supports 500K bit/sec downstream and 120K bit/sec upstream. ■

Technology update

Wireless technology puts users in place

■ BY PRADEEP IYER

As companies roll out wireless networks, one area of concern is how to automatically segment wireless users into the correct virtual LANs already established on the wired side. VLAN membership on wired networks typically is defined by the physical Layer 2 switch or Layer 3 router port to which a user is connected. But with wireless, users aren't tied to a physical port.

To address this problem, advances in wireless authentication have led to role-based VLAN association. This method of automatically deriving the correct VLAN membership uses a number of standard authentication methods, such as HTTP-based captive portals and 802.1X, which has become the authentication mechanism of choice.

Consider this scenario. Wireless users in a finance department might be connected securely to the Finance VLAN using a secure-link encryption method such as Wi-Fi Protected Access. However, once they roam to another access point, they no longer necessarily have access to the Finance VLAN and can't use their network resources. Reconfiguring the network to make each VLAN accessible from every point across the entire company is not a viable solution.

However, 802.1X port-based authentication provides a framework for authorizing station access to Ethernet and wireless LANs. 802.1X uses Extensible Authentication Protocol (EAP) to relay port-access requests between LAN stations (supplicants), Ethernet switches or wireless access points (authenticators), and RADIUS servers (authentication servers).

■ HOW IT WORKS

Role-based VLANs

Network administrators can segment wireless users on a network with role-based VLANs that use 802.1X authentication.

Finance user

Access point   WLAN switch   RADIUS server

Marketing user

① After users connect to the access point, the WLAN switch sends a request back to the users for ID.

② The users respond with IDs, and the WLAN switch forwards them to the RADIUS server.

③ The RADIUS server accepts the access requests. The switch extracts user-identity information to automatically place the users in the correct VLANs. An acceptance message is sent to the users.

④ Once the users are placed into the role-based VLANs, all subsequent traffic is forwarded to the correct VLAN, such as Finance or Marketing.

The central mechanism used to protect users in Wi-Fi networks is based on data encryption and user authentication — not typically by roles derived from an authentication method. Role-based VLAN association with 802.1X is attractive because it provides logical segmentation of workgroup traffic, and easier integration with security and traffic-engineering policies configured on wired networks.

Network administrators want to keep the same Extended Service Set IDs (ESSID) and encryption profiles for all users, and assign users in different workgroups to different VLANs as they enter the wireless LAN (WLAN), based on attributes already configured on the authentication server. Without role-based VLANs, this isn't possible unless you make a lot of changes to WLAN configuration by introducing new ESSIDs for each user group. This represents a significant capital investment and operational expense.

A WLAN switch can support a variety of user roles with different access rights and VLAN associations. It also can support a variety of server rules from which to derive a user role, such as the RADIUS attributes in the access-accept message from the RADIUS server. For example, a server rule can be defined to extract the value of a specific RADIUS attribute (say Attribute 11, Filter-Id) and use the value as the role. In 802.1X authentication, the client authenticates to the RADIUS server through a WLAN switch. The WLAN associates a VLAN to the client based on the role derived by applying the server rules.

The WLAN switch puts the client in unauthorized state once 802.11 association with an access point is complete. In this state, only 802.1X EAP packets generated by the client are forwarded through the WLAN switch. The WLAN switch sends an EAP Request-ID, a user identity request message, to the client. The client responds with an EAP Response-ID message. The WLAN switch encapsulates the EAP Response-ID as a RADIUS access-request message and forwards it to the RADIUS server.

If authentication is successful, the RADIUS server sends an access-accept message to the WLAN switch. This message identifies different user attributes such as role and access rights. The WLAN switch then parses this response to determine into which VLAN the client should be placed.

Using this information, the WLAN switch places the client in an authorized state and sends an EAP Success message. It then forwards all future data traffic from the client to the right VLAN. Upon receiving the EAP Success message, the client starts a Dynamic Host Configuration Protocol transaction to get an IP address on the role-based VLAN.

Iyer is a principal software developer at Aruba Wireless Networks. He can be reached at piyer@arubanetworks.com.

Ask Dr. Internet By Steve Blass

What is the Internet Storm Center? How can we contribute firewall logs to its network security monitoring process?

The Internet Storm Center grew out of the SANS Institute's Consensus Incident Database project, which began monitoring global Internet traffic in November 2000. The center is a free service that collects intrusion-detection information from the Internet to identify new attacks and provide information about the types of attacks being mounted against Internet resources. At www.incidents.org you can find lists of the top 10 ports being attacked, ongoing attack trends and other security information, including an online ticker showing current trends. To submit your traffic logs to the system, go to www.dshield.org/howto.php. Automated client software is provided for several Windows and Linux firewalls. The client software automates the process of finding the appropriate portion of your firewall logs and e-mailing the information to Dshield. Logs also can be submitted through a Web form interface or directly through some firewall hardware equipment. Registering with Dshield is encouraged but not required to participate in the reporting service.

Blass is a network architect at Change@Work in Houston. He can be reached at dr.internet@changeatwork.com.

Dr. Jim Metzler

Dr. Jim Metzler is widely recognized as an authority on both network technology and its business applications. In over 28 years of professional experience, Jim has assisted tens of vendors refine their product strategies and simultaneously helped over a hundred enterprises evolve their network infrastructure.

Managed Services

> **What is a managed security services provider (MSSP)?**

An MSSP is a company in the business of providing electronic security services on a third-party basis. (MSSPs such as Unisys are preparing to support physical security such as implementations of Lenel access control systems or IP-enabled video surveillance cameras that might be remotely monitored from a security operations center.) There are a wide range of MSSPs—from companies that provide one or two very specific security services to companies that provide a large number of security services.

> **"The principle benefit of using an MSSP is that it gives a company access to skilled resources."**

> **What are the benefits of outsourcing your enterprise security infrastructure to an MSSP?**

Driven either by concerns about their own financial vulnerability or possibly by government regulation, security is one of the top issues in virtually every IT organization. Moreover, security is also very visible—if a company has a security breach, it is often well-publicized both within and without the company. The principle benefit of using an MSSP is that it gives a company access to skilled resources.

> **Does a company relinquish control of security services when using an MSSP?**

That is the key risk associated with any sort of outsourcing relationship. Given that, it is important to turn this concern into the key criteria that a company uses when choosing what they will outsource, as well as to whom they will outsource. In particular, when a company is considering using an MSSP, the company needs to ensure that the MSSP has processes that are flexible enough to ensure that the company is not giving up an unacceptable amount of control.

> **Which strategic security functions should remain in-house and which can be outsourced?**

There is one function that absolutely must be outsourced, and that is doing a security audit. Each company must have regular security audits performed. The MSSP that is doing the audit must of course have expertise in this area and must also be clear of conflict of interest—it cannot be an organization that is providing any other security functionality for the company.

As a general rule, companies that fit the following criteria should outsource functions:
• The company is not good at performing that functionality, nor do they foresee developing the requisite expertise;
• The company feels confident that they have found an MSSP with the expertise;
• The company feels confident that it can maintain a high level of control while using the MSSP.

> **What should a company look for in a potential MSSP?**

A good security approach should embrace a multilayered security infrastructure that requires multiple technologies, process and procedures to be breached. In evaluating security outsourcing, an organization should make sure that they do the following:
• Obtain clear and concise Service Level Agreements
• Clearly understand the roles and responsibilities of both the outsourcing provider and the in-house staff.
• Come to clear agreement on security incident severity levels and the desired actions should an incident occur.
• Look for up-to-date accreditation or certifications for the personnel who work for the MSSP.
• Examine what security tools are used by the MSSP, particularly any tools that the MSSP might have developed that provide important functionality that the company currently lacks.
• Are the MSSP's processes thorough and detailed enough to show a thought-out, well-documented approach to providing security?
• Choose an MSSP that has a strong, demonstrable track record of providing the security services of interest to a wide range of customers, including some that are similar to the company in question.
• Check to make sure that the MSSP can provide the services in the local geography or required language.

Dynamic DNS zeroes in on IP addresses

O ur sister-in-law Lydia runs a pre-school and wanted to have a wire-less Webcam so parents could see their little darlings. But first there was the problem of her DSL connection. One of the more annoying issues with consumer broadband links is the common use of dynamic IP addresses. You get a different address each time you restart your connection.

While this is generally acceptable, it is a problem when you need to run a server — you can hardly serve anything if you can't be found. Let's say you are setting up a branch office, and you have a dozen people who want to pick up e-mail. Rather than having them all go out over a shared DSL connection you might decide to give them an in-house e-mail server that, because it is central-ized, also lets you archive incoming messages.

So if the IP address is not guaranteed to be the same twice, what can you do? Well,

dear reader, there is a simple answer: dynamic DNS. Essentially this is a DNS server that can be updated frequently. At some computer at the branch location you run software called an update client that periodically talks to the dynamic DNS server and reports the current IP address. Voilà! The branch office now can be found.

There are many dynamic DNS servers on the 'Net. One we have used is DynDNS, owned by Dynamic DNS Network Ser-vices. DynDNS provides a range of charged-for services, but it also offers free entries for up to five host names per person.

A name mapped to a dynamic IP address will be suffixed by one of more than 40 domains that DynDNS makes available — for example, spokane-office.gotdns.com.

We used an update client called DynSite to make our sister-in-law's preschool Webcam — a D-Link Systems DCS-100W Air 2.4-GHz wireless network Internet camera — accessible to parents through DynDNS. (You can read more about DynSite at www.nwfusion.com, Doc-Finder: 7929; more about the D-Link cam-era is available at DocFinder: 7928)

The parents love it; one of them said

she regularly gets to eat lunch with her child because of it. You can see D-Link's write-up of the school's use of its gear (further information at DocFinder: 7929).

Getting the picture

The camera connects over 802.11b, or 10/100M bit/sec, Ethernet and can pro-duce VGA video at up to 20 frames per second. Even when we were 30 feet from the wireless base station, which is sur-rounded by metal filing cabinets and with a wall in the way, we could reliably deliver streaming video across the Internet at around one frame per second.

The DCS-1000W has both automatic and manual modes for gaining control, expo-sure and white balance; and shutter speeds from 1/60 second to about 1/15000 second. It also will function down to an illumination level of 2.5 lux at f1.4.

The camera has a built-in Web server, and through the utility software you can set the frame rate and image size; and define users, their passwords and access levels. The actual imaging at the receiving end can be done via a Java applet or an ActiveX control. So far we have been unable to get the ActiveX component to work outside the LAN, but the Java applet

works just fine across the Internet.

This camera has many more sophisti-cated features (all sorts of event triggers, including motion in the field of view and external devices) that make it a very good monitoring solution. The bundled soft-ware lets you monitor up to 16 cameras simultaneously. You also can record video to the hard-disk drive.

It was pretty easy to get running and configured, although as with all of these wireless products, the average non-tech-nical user would have a tough time. The only complaint we have is that we can't "un-frame" the Java applet from the default Web page the camera generates. If we try to load the applet directly from the camera we can't authenticate, and therefore we can't receive the video stream. The problem with the default Web page is that it is techie — it has but-tons to switch triggers on and off, which isn't relevant to what the parents want. C'est la vie. D-Link is, we hope, finding a fix for us.

All in all, a great product, and if we can get a fix for accessing the camera without the default Web page, it might even achieve a terrific rating.

Rate us at gearhead at gibbs.com.

Cool Tools

**Quick takes
on high-tech toys
By Keith Shaw**

Palm, Sony launch new handhelds

P alm and Sony each launched new PDAs last week, with features such as a landscape display option, faster processors and additional memory.

Palm's new handhelds include two models in the Tungsten line (the Tungsten T3 and Tungsten E) and the $99 Zire 21, the successor to last year's budget-model Zire.

The Tungsten T3 ($399) includes a high-resolution color screen (320-by-480-pixel TFT) that can display in land-scape and portrait modes. When used in landscape for-mat, the handheld display offers better viewing for appli-cations such as spreadsheets, movie clips and Web page browsing. The screen on the Tungsten T3 is about 50%

larger than any other Palm-branded device, Palm says.

Other features include built-in Bluetooth connectivity, 64M bytes of RAM (with 52M bytes available for users), a 400-MHz XScale processor and an expansion slot that supports and MultiMediaCard (MMC) and secure digi-tal media.

The $199 Tungsten E device includes 32M bytes of RAM (with 28.3M bytes available for users), a 320-by-320-pixel color display and an improved five-way navigation button for easier one-handed use. It runs on a Texas Instruments OMAP 311 ARM processor and includes multi-media software that lets users listen to audio files, watch movie clips or view photos. The device also has an expansion slot that supports secure digital and MMC media.

The Zire 21 comes with 8M bytes of memory (four times as much as the original Zire, with 7.2M bytes available for users), a faster processor (the 126-MHz Texas Instruments OMAP processor), and the newest versions of the date book, address book, notepad and to-do-list applications.

Palm also updated its core applica-tions for the new Tungsten handhelds. New features include an agenda view (shows appointments, daily tasks and a new "year view"); scheduling that lets

**The Sony Clie PEG-TJ35 includes a
four-way navigation button.**

users beam multiple appointments with a single command; additional contact information (including multi-ple contact addresses, more spaces for phone numbers and e-mail addresses, instant-messaging names, Web site URLs and birthdays); a larger memos and notes field; and more built-in Microsoft Outlook compatibility.

New from Sony

Sony's latest Clie handhelds, the PEG-TJ25 and PEG-TJ35, are expected to be available by the end of the month, and priced at $200 and $250, respectively.

The devices include a high-resolution (320-by-320-pixel) color LCD screen and an enhanced Jog Dial navi-gator that now provides four-directional movement for quick access to applications, Sony says. Two applications, Clie Memo and Clie Viewer, have been pre-installed onto the read-only memory of the new devices.

The handhelds include a 200-MHz processor, integrated MP3 player, Decuma Latin's handwriting-recognition soft-ware and the Picsel Viewer application for viewing native Word, Excel and PowerPoint files. The TJ35 model includes 32M bytes of RAM (23M available for users), and the TJ25 includes 16M bytes of RAM (11M bytes for users). Both devices include a Memory Stick Pro expansion media slot for memory expansion up to 1G byte, Sony says.

**Palm's Tungsten T3 includes a landscape display
mode, making it easier to view spreadsheets.**

Shaw can be reached at kshaw@nww.com.

Screw security.
It's uptime that counts.

Face it...worms, viruses, unexpected traffic surges, they're going to get you. And, unfortunately, security systems don't identify problems...until the damage is done. And as we all know, it's impossible to stay functional as your network is slowly grinding to a halt or worse...shutting down entirely. So the real issue is network uptime. Imagine a network system so intelligent, it can quickly identify the difference between good traffic and bad, with the sophistication to immediately throttle down and control specific streams of traffic, while allowing others to enter and flow freely. It's a new way of dealing with a very old problem...maintaining performance and keeping your network up.

captus
N E T W O R K S

Beyond Intrusion Prevention: Predictable Network Availablity
For a quick, free demo...visit us at **captusnetworks**.com

opinions!

EDITORIAL
John Dix

Sun fights back with innovation

un's announcement last week that it will post a larger-than-expected loss in its fiscal first quarter is the latest evidence that Sun is trying to exorcise demons.

With increasingly powerful Intel-powered boxes running Linux, which call into question the need for Solaris on Sparc, Sun has only managed to post three profitable quarters — slim at that — in the last eight.

But Sun is fighting back, touting new low-cost products and an interesting new software story.

Regarding the former, Mark Tolliver, executive vice president of marketing and strategy and chief strategy officer, says Sun already offers some Intel-based servers for less than Dell, and much less than HP and IBM. "And our one- and two-way entry-level Sparc servers are 65% cheaper than they were two years ago," he says.

But can Sun ramp up the low-cost business fast enough to offset declining demand for higher margin products? "We have to do it," Tolliver says.

To help round out the balance sheet, Sun also is looking to drive up sales of higher margin software. The company led the recent SunNetwork user conference with a software announcement, the Sun Java Enterprise System (formerly Project Orion).

As we reported, this is an infrastructure software suite that includes homegrown code and products acquired over the past five years. It includes identity management controls, Web and application servers, messaging, calendaring and portal tools, and clustering services for availability.

Delivering the tools as an integrated suite ensures the pieces work together. And Sun will update the suite on a fixed, quarterly basis and deliver a single installer for automated updating, which saves customers the hassle of keeping up with and keeping track of various versions and releases.

The kicker: The entire suite costs $100 per employee, per year, which includes service, support and a certain amount of professional services. So no more complex negotiations, no huge upfront cost, no mixed billing cycles and no more auditing.

Tolliver, who used to run Sun's software business, says he thinks this changes the software game. With current software licensing practices, "once you write me that first $6 million check, I have an almost irresistible urge to seek the next one," he says.

The pricing scheme and delivery model are indeed compelling, but the question is whether the software itself is compelling enough to attract customers. As attractive as it looks, most companies already are committed to products in these categories.

— John Dix
Editor in chief
jdix@nww.com

Source verification assumption

Regarding Mark Gibbs' Backspin column "Running the numbers on source verification" (www.nwfusion.com, DocFinder: 7924): One problem with source verification is that it assumes the people e-mailing you want you to get their message badly enough to go through the second response. Probably not a good idea for your main sales or customer service/support addresses.

Source verification also assumes that the sender checks mail as (or more) often than the time frame you need to receive the initial message. You'd hate to miss that hot tip sent Friday afternoon because the sender didn't check e-mail again until Monday.

Finally, source verification assumes the sender is going to recognize the "from" address or subject and not decide the challenge is spam and just delete it.

Cory Jaeger
Network manager
D.C. Everest area schools
Weston, Wis.

Punishing the innocent

Regarding Johna Till Johnson's column "MCI charges: Oklahoma gets it right" (DocFinder: 7925): The feds really missed the boat. The only people being punished by their order to no longer book new business with MCI are the workers who deliver the services. These people — working-class, tax-paying individuals — had no authority to make decisions relative to the fraudulent finances. The people left at MCI are being punished a second time, as most of them saw their retirements disappear because of Bernie Ebbers and his cronies, and now our government wants to ensure that they get punished again. Punish the guilty and leave the innocent

E-mail letters to jdix@nww.com or send them to John Dix, editor in chief, Network World, 118 Turnpike Road, Southborough, MA 01772. Please include phone number and address for verification.

alone; do business with MCI and put Bernie in jail.

Pitt Kays
Marshall, Mo.

Business ethics and politics

In his letter to the editor regarding the MCI affair (DocFinder: 7521), David Easter blames Republicans for our nation's economic woes and for permitting businesses to continue to be "rotten to the core." Here are some facts he should consider.

The scandal with Enron did not begin after Sept. 11, 2001. The key piece of the government's evidence is a memo dated Feb. 6, 2001, only 17 days after George W. Bush became president. That memo implies years of accounting improprieties. One analyst was warning his listeners about the problems with Enron's "value" as far back as 1997. Now, if I recall, our president through the last portion of the 1990s was Bill Clinton, a Democrat, and that is when these scandals had their origin. The problems with companies such as WorldCom/MCI, Tyco and Global Crossing also started in the late 1990s.
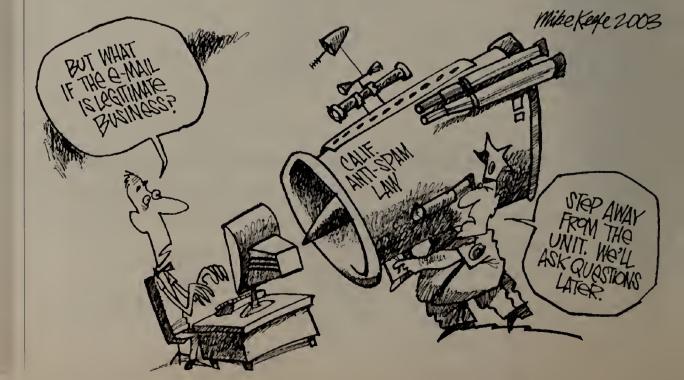
As for recessions being inevitable with Republican presidents, it was not Bush who signed the free trade agreements that have permitted American businesses to export American jobs to third-world countries; it was Clinton.

It is an error to place the blame for our businesses' lack of ethics, and our woes as a result, on either political party. The blame lies squarely with those "businesspersons" whose greed compel them to steal from their employees and propels them to destroy the businesses they are supposed to be managing to feed their own appetites. This pillaging of American workers will continue regardless of who is president until American investors, consumers and workers stop letting them do it by refusing to work for them, buy from them or invest in their scams.

Dale Greenlee
Independence, Ore.

INTRANET ADVISER
Daniel Blum

Federated ID gains momentum

Early adopter pioneering challenges notwithstanding, federated identity is worth the effort.

Gone are the days when Midwestern manufacturing companies had to roll their own XML security to achieve single sign-on across Webbed supply chains. A spate of product announcements from vendors such as IBM, Oblix and RSA Security are bringing increased levels of turnkey Security Assertion Markup Language support.

Federation — the practice of authentication and identity information exchange across different security or technology domains — is spreading through different vertical industries. Like identity management overall, the drivers behind federation are a need for reduced sign-on, application integration and regulatory compliance. Early adopters report that even though you have to pay to play during the early stages of federated identity deployment, ROI is there for the taking. A Shibboleth project leader at The Pennsylvania State University who implemented SAML in 2002 reports an 85% drop in help desk calls.

Financial services companies have been using identity networks such as SecuritiesHub for years. There's also been significant activity in industries such as mobile telecommunications, insurance, automotive, aerospace, manufacturing, government, travel and higher education.

Organizations are adopting federated identity for many use cases. Some require basic SAML authentication assertions for SSO in business-to-business scenarios. Others are looking at Liberty-Alliance-enabled products for consumer accounts linking with e-business affiliate partners. Others need SAML and/or Liberty for SSO with benefits suppliers, outsource partners or internal applications. While use cases today focus on browser-based SSO, federated identity also will become part of Web services deployments, providing back-end process integration and transactions in the longer term.

In an exciting twist, many business managers see more than just ROI; they see federated identity enabling competitive advantage. Demand is coming bottom up from business units and top down from IT infrastructure organizations.

But expect a few hardships. Many use cases require technically complex user name mapping or attribute information, and you have to specify and test the way the protocols will operate. Some vendor products are not as interoperable as their marketing literature would imply. But the most difficult issues are non-technical, such as getting executive buy-in, establishing agreements with partners, or passing legal reviews and risk assessments for this new way of doing business.

Early adopter pioneering challenges notwithstanding, federated identity is worth the effort. And fortunately, today's products are increasingly functional, and with vendors such as BEA Systems, IBM, SAP and eventually Microsoft jumping aboard the SAML train, there's a realistic expectation that stronger, identity-based security mechanisms will become native to platforms, tools and today's add-on security middleware. I haven't seen anything this exciting in this area since Multi-purpose Internet Mail Extensions began spreading like a prairie fire across the world of Internet mail in the early 1990s.

Blum is senior vice president and research director with Burton Group, an integrated research, consulting and advisory service. He can be reached at djb-feedback@earthlink.com.

CACHE ADVANCE
Linda Musthaler

HP should stay in the PC business

HP leaving the PC business would be bad for the industry and a real loss for customers.

Eighteen months after the largest merger ever in the technology industry, some analysts have grown impatient waiting to see if HP will reign supreme in the PC market. The thought is, if HP can't compete against Dell, it should pull out of the PC business, in effect ceding the market to Dell.

Joseph Beaulieu, computer analyst with the financial advisement company Morningstar, doesn't mince words, saying, "It would probably in the long run enhance [HP's] profitability to get out of [the PC business]."

Even Mike Elgan, a spokesman for the HP user group Interex, boldly proclaims, "HP should get out of the low-cost, low-margin, low-innovation PC business and focus all its energies on product and service areas where HP's technological superiority matters."

The sting of such comments prompted HP to issue a July press release in which Jim McDonnell, a vice president in the Personal Systems Group, said: "We're in the battle for the long term, we're providing great products at aggressive prices to our customers, and we look forward to continuing to build on our momentum in the second half of the year."

I'm glad to see HP reaffirm its commitment to the market. HP leaving the PC business would be bad for the industry and a real loss for customers. Why? I can sum it up in four words: innovation, competition, simplicity and choice.

When it comes to developing new products for the commercial and consumer sectors, HP is an innovator. In 2002, HP was awarded 1,385 U.S. patents, making it No. 9 on the list of top companies receiving patents from the U.S. Patent and Trademark Office. Not all of these patented technologies made their way into HP's PC products, of course, but the fact remains that HP spends a significant amount of money on research and development for its PC products, especially when compared with Dell. If HP exits this market, we'll lose some very innovative thinking.

Then there's the competition factor. IDC reminds us quarterly of the tight race for PC market share. When HP recently announced a slight rise in its worldwide market share, Dell responded by lowering prices. If HP left the PC market, Dell wouldn't feel the pressure to continuously lower prices, and customers would have little leverage to get Dell to strike a bargain. Such competition is healthy for the market, forcing the players to become more efficient to stay in the game.

Next, there's simplicity. Many companies like to limit the vendors they deal with, especially when it comes to computer technology. For those who get their enterprise systems from HP, they can streamline their purchasing and support by buying HP PC products, too.
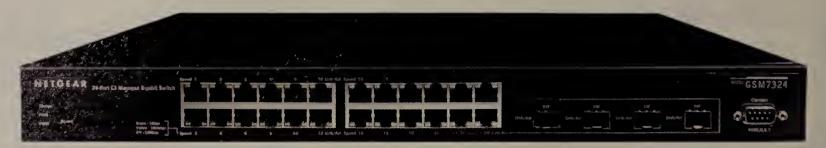
Finally, customer choice is still important. The sad fact is that healthy competition in the PC market is dwindling. The market has already consolidated too much, becoming basically a four-horse race (Dell, HP, IBM and "white box" or reseller brand). I miss seeing companies such as AST Research, Leading Edge and Wyse in the PC market. And I'm thrilled that Gateway just signed a deal to provide the U.S. Defense Logistics Agency with tens of thousands of new computers over the next few years. Gateway has been perilously close to being forced out of the PC market, and this new contract might be the boost it needs to stay in.

So HP's Personal Systems Group has suffered through a few dismal quarters. That's unfortunate, but not reason enough to quit the market. Maybe the company just needs to reevaluate its game plan, as IBM did a few years ago when it dropped out of the consumer PC market. Instead, IBM's PC Division found its niche and is thriving today with commercial customers that buy IBM enterprise systems and services.

Beaulieu and Elgan are looking at HP's bottom line and not necessarily what's good for customers. I'd be very sad if HP gave up trying to compete profitably in the PC market. Then we might be hearing, "Dude, yer getting a Dell … because there's nothing else to choose from."

Musthaler is vice president of Currid & Company, a Houston technology assessment firm. She can be reached at linda@currid.com.

Face-off

Is the GPL good for the software industry?

A free software advocate and a SCO Group exec debate the merits of the GPL.

Yes, by Bradley Kuhn

The GNU General Public License has a positive effect on the software industry. Vibrant software sharing defended by the legal protections of GPL inspires growth and advancement, just as publishing and sharing research results invigorates fields such as physics, mathematics and psychology.

Software advances through incremental improvement. Paradigms shift, standards change, and methods are reinvented. Access to others' results, both successes and failures, spurs rapid growth. The widespread adoption of Internet technology started from widely shared and incrementally improved free software. That process continues today.

The GPL does change the ethical implications of our development activity. Sharing software is encouraged; sharing improved versions of software is rewarded. The primary industrial mechanisms and business models for software — support, contracted customization and improvements for hire — thrive and are equalized under this model of freedom. Software is now, as it really always has been, a service rather than a product. The GPL lays the ground rules, ensuring that no particular developer or company holds power over any other, and no one controls the software users.

Users face a free marketplace. A vendor who distributes under GPL does not lock you in to their product. If you don't like your contractors, you fire them and hire new ones. You have the source code, and the means and rights to modify it, so you can do the work in-house. Software companies can be held accountable by their customers and must actually show the value of the expertise that they add to the software.

Most industries that are primarily intellectual in nature, such as software, law and auto mechanics, thrive best as a market for experts. Information about the field is publicly available, taught in universities and swapped among practitioners. But experts who can leverage their knowledge into clear results for clients move to the top. Activity of experts under the umbrella of GPL forms a meritocracy, and yields a shared commons that profits all.

Admittedly, some business models don't function in that meritocracy. The model whereby you bamboozle the world into running your proprietary software and extract an exclusionary licensing fee from each individual — who cannot fix bugs, make improvements or adaptations or get support services from anyone but you — has already begun to collapse. Trade-secret proprietary software, based on keeping knowledge away from users and programmers, now fails the test of business effectiveness as well as ethical propriety. A new IT economy, properly based on software engineering's scientific roots, has emerged. The GPL underpins that new economy, as it puts users, developers, customers and academics on equal footing to improve and profit from the shared commons.

Success in this new industry will not be determined by exclusionary licensing deals, but by the ability of your software engineers to understand and improve the commons. The GPL creates a fair and competitive software industry that functions as a scientific endeavor, not a snake oil sale.

Kuhn is executive director of the Free Software Foundation, a charitable organization in Boston. He can be reached at bkuhn@fsf.org.

No, by Chris Sontag

The General Public License is not good for the software industry for a variety of reasons. These include:

● The GPL is full of contradictions and could be interpreted in a number of different ways. At the outset, it bears noting that the GPL is as much a political manifesto as it is a quasi-legal document. It is replete with ambiguities. More importantly, it has never been tested in court. Nobody knows what judges or juries will do when presented with the GPL and asked to find that someone violated its terms, or when someone tries to use it as a defense to a copyright or patent infringement claim.

● The GPL's authors have one point in mind: to destroy the value of proprietary software. Richard Stallman, founder of the free software movement, was recently quoted as saying: "Proprietary software is antisocial and shouldn't exist." If the authors of the GPL have their wish, there will be no cost to software. Why is that so bad for end users? If software is free, companies won't be around to service that software, provide customer support and produce upgrades. How will software companies afford to pay for salaries, benefits and keep people gainfully employed? It won't happen with free software.

● IT companies avoid porting to software that is licensed under the GPL. SCO recently received a letter from a company supporting our current legal battles, stating: "We have resisted porting our software tool to Linux because of the fear of seeing our source code published on the Internet shortly thereafter." We believe that for Linux to flourish in the future, a license other than the GPL will have to be prescribed.

● The GPL causes software innovation to stagnate. In this same letter that SCO received, the writer states, "Small developers like us used to be the lifeblood of the computer business — innovating and bringing fresh ideas and products to the marketplace. How can this continue if we are supposed to donate all of our efforts?" The GPL essentially prohibits a company from taking a software product like Linux, writing proprietary applications and add-ons, and then selling that software without showing anyone what was done to it.

● The authors of the GPL wrote the license in such a way that it would govern the use, distribution and copying of software that was licensed under the GPL. These are the same items governed by the U.S. Copyright Act. The Copyright Act pre-empts any claims that are governed regarding use, distribution and copying. Because of this, SCO believes the GPL is pre-empted by federal copyright law.

SCO believes that there are better licensing models available which, unlike the GPL, are not in conflict with U.S. copyright law. These licenses give developers greater incentives to innovate without destroying the value of proprietary software. Until the legality of the GPL is fully tested, organizations that rely on open source software released under the GPL will continue to take an unnecessary risk. The only way in which this risk can be mitigated is for the GPL to change, or for developers to work under more flexible licenses.

Sontag is senior vice president and general manager of the SCO source division of The SCO Group. He can be reached at csontag@sco.com.

More online!

Log on to Network World Fusion to voice your opinion. Face-off authors Bradley Kuhn and Chris Sontag will add their thoughts to the discussion.

DocFinder: 7922

Performance:
Superior power and reliability.

Thirty years ago Ethernet was just an idea. Today it's the technology that drives your business. For nearly three decades, 3Com® has delivered the high-performance Ethernet switching that keeps networks running reliably, and provides the tools for people to work productively and businesses to compete successfully. For proven performance and value, choose 3Com 10/100 switches.

Enterprise-class features and flexibility
Feature-rich 3Com SuperStack® 3 Switch 4400 family switches are ideal for high-performance environments running IP telephony, Power over Ethernet, security, or where a highly optimized, resilient network is needed.

Small office affordability and convenience
Economical 3Com SuperStack 3 Switch 4200 family switches offer Layer 2 switching that's easy to install, operate and afford. Built-in stacking, wire-speed performance, ease-of-use features and standards-based operations help ensure that you'll be up and running in no time.

Get CASH BACK on every 3Com SuperStack 3 Switch 4400 and Switch 4200 you buy... plus the chance to win a classic 1973 Chevrolet Corvette!*

SuperStack 3 Switch 4400 Family

SuperStack 3 Switch 4200 Family

NEW!
▶ Price reduced— Save up to $220!
▶ Next business day advance hardware replacement.

3COM
Possible made practical™

73 VETTE

Go to www.3com.com/73corvette/promo27 for your chance to win!

Plus, get   **$60 CASH BACK** for every Switch 4400 you buy
                  $50 CASH BACK for every Switch 4200 you buy
                  $150 CASH BACK BONUS when you buy any 3

HURRY—OFFER ENDS NOVEMBER 30, 2003!

1973
30 YEARS of Ethernet
2003

NetworkWorld
Review

WAN MONITORING TOOLS
Visual UpTime edges Concord's eHealth in test of six software packages

■ BY BARRY NANCE, NETWORK WORLD GLOBAL TEST ALLIANCE

Your company's transactions, queries, documents, intranet data and files are its lifeblood, and your network's connections are the arteries that carry that blood. Keeping those connections healthy is more than just prudent. It's critical. No company wants to see its network in intensive care — or the morgue.

Many vendors offer monitoring software, devices or combinations of both to help you maintain WAN links at the peak level. These vendors promise their tools will alert you when outages occur, pinpoint the root cause of the outage and help you reestablish communications immediately. They claim to produce useful reports showing utilization trends, outage statistics, service-level agreement (SLA) compliance and other information. Vendors say the tools are easy to use, scale well, integrate with network management systems, handle any and all protocols, and have lots of additional features, such as the ability to prioritize network data based on quality-of-service parameters you provide.

To find the best WAN monitoring tool for your network, we invited vendors to submit their products to our lab. We tested Network Instruments' Observer 8.3 software and rack-mountable WAN Probe with a pair of T-1/E-1 analyzer taps; Neon Software's CyberGauge 5.0 software; Adtran's IQ 710 with traffic shaping and N-Form 1.4 monitoring software; Visual Networks' Visual UpTime 7.1 and Analysis Service Elements (ASE — DSU/CSUs augmented with link monitoring capabilities); Concord's eHealth 5.6 software; and Allot's WiseWAN 401 Network Application Priority Switches (link monitoring, shaping and controlling devices) and WiseWAN Network Application 5.2 Enterprise software.

Visual UpTime was the best product for keeping WAN links up and running smoothly and wins our World Class award. Although it only works with Visual Networks' DSU/CSU devices, Visual Up-Time's precise and accurate monitoring ability is unsurpassed. Its many reports are practical and well designed, the user interface is intuitive and responsive, and it scales well.

For heterogeneous networks, Concord's eHealth is a World Class winner for its superior reports and amazing breadth of recognized and supported devices.

All the products did well in our tests. They proved themselves worthy, reliable tools for monitoring critical WAN links.

Hardware vendors leverage their sales by bundling or offering software that

Net Results

RATING 4.6 WORLD CLASS WINNER
Visual UpTime 7.1 software and ASEs (DSU/CSUs)
Company: Visual Networks, (301) 296-2300, www.visualnetworks.com **Price:** Starts at $1,000 per site. **Pros:** Accurate, precise WAN monitoring tool. **Cons:** Only monitors Visual Networks devices.

RATING 4.5 WORLD CLASS WINNER
eHealth 5.6 software
Company: Concord Communications, (800) 851-8725, www.concord.com **Price:** Typical licenses range from $100,000 to $150,000, depending on infrastructure size. **Pros:** Excellent vendor-neutral data gathering and reporting tool. **Cons:** Network discovery and baselining can take a while.

RATING 4.3
IQ 710 DSU/CSUs and N-Form 1.4 software
Company: Adtran, (800) 923-8726, www.adtran.com **Price:** An IQ 710 with traffic shaping costs $2,395. N-Form software costs $6,000. **Pros:** Good traffic-shaping ability. **Cons:** No printed documentation; user interface isn't as intuitive as that of the other products.

RATING 4.1
Observer 8.3 software and WAN Probe
Company: Network Instruments, (952) 932-9899, www.networkinstruments.com **Price:** $2,895 for Observer, $6,000 to $15,000 for each probe. **Pros:** The perfect tool if you sometimes need to drill down from a top-level summary to individual problem packets. **Cons:** Offers fewer reports than Visual UpTime or eHealth.

RATING 3.9
WiseWAN 401 hardware and WiseWAN Network Application 5.2 Enterprise software
Company: Allot Communications, (408) 988-8100, www.net-reality.com **Price:** From $6,000 to $32,000. Includes monitoring devices, software and relational database. **Pros:** Well-designed user interface; supports traffic shaping. **Cons:** Only useful with WiseWAN hardware.

RATING 3.7
CyberGauge 5.0 software
Company: Neon Software, (925) 283-9771, www.neon.com **Price:** Starts at $295 for five monitored devices. **Pros:** Simple, easy-to-operate link monitor. **Cons:** Suitable only for small networks.

The breakdown	Visual UpTime 7.1 WORLD CLASS Winner	eHealth 5.6 WORLD CLASS Winner	IQ 710 & N-Form 1.4	Observer 8.3 & WAN Probe	WiseWAN 401 & Network Application 5.2 Enterprise	CyberGauge 5.0
Reports 20%	5	5	4	4	4	3
Discovery 20%	3	5	5	3	3	4
Ease of use 20%	5	3	4	4	4	3
Protocols 10%	5	5	5	5	5	5
Network management integration 10%	5	5	5	5	5	5
Documentation 10%	5	5	3	5	3	3
Installation 10%	5	4	4	4	4	4
TOTAL SCORE	4.6	4.5	4.3	4.1	3.9	3.7

■ **Scoring Key: 5:** Exceptional; **4:** Very good; **3:** Average; **2:** Below average; **1:** Consistently subpar

WAN tools

works only with their devices. In contrast, software vendors work hard to support as many devices as possible. This can pose a dilemma for companies planning to expand or upgrade an existing network.

Not surprisingly, we saw the best and most-detailed monitoring of our WAN links from products that merged a vendor's software with its hardware devices. Visual UpTime gathered statistics from and sent control commands to Visual Networks' own DSU/CSUs; Adtran's N-Form software provided WAN monitoring for the company's IQ 710 DSU/CSUs; and the WiseWAN Network Application Enterprise software worked with the WiseWAN 401 monitoring and traffic-shaping devices. Similarly, Network Instruments' Observer WAN monitoring relied on the presence of a WAN Probe located at the other end of a monitored link.

On the other hand, Concord's eHealth and Neon Software's CyberGauge gave us support for a range of network devices, but didn't monitor as closely nor deliver the level of detail that, for example, Visual UpTime did.
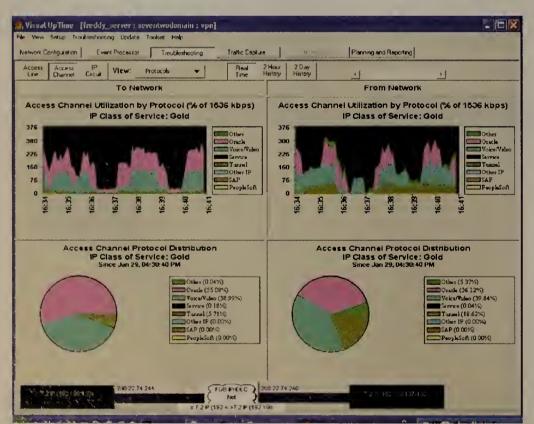
Discovering and reporting problems

The key to Visual UpTime's success is its close relationship with its ASE devices. The ASEs continually measure link availability and activity on a second-by-second basis

tual circuit (PVC). We even found that for the sake of accuracy, we could exclude scheduled maintenance periods from Visual UpTime's calculations of uptime and bandwidth utilization. UpTime used the data from the ASEs to clearly show us outages and traffic levels. It also showed us several frame relay metrics, such as per-port and per-PVC throughput, overall utilization, by-protocol utilization, bursting above the committed information rate (CIR) and network congestion identified by the presence of frame relay internal throttling mechanism packets.

Combining Adtran's IQ 710 traffic-shaping DSU/CSUs and N-Form software not only monitors links for availability, but also recognizes application-specific traffic and prioritizes that traffic during busy periods. It can identify more than 300 kinds of application-level network datastreams, including Citrix WinFrame, HTTP, AOL Instant Messenger and Napster messages. Both the IQ 710s and N-Form track and display the same frame relay metrics as Visual UpTime, although with not quite much fine detail.
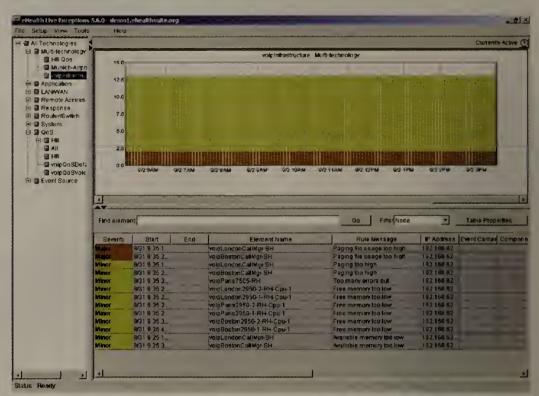
While some Allot WiseWAN models include DSU/CSU functionality, the model 401 units we tested were pure monitoring devices. An Ethernet tap connected the 401 to the link between the local network and router. Allot also offers a WiseWAN unit for monitoring broadband (DSL) connections. The WiseWAN WANXplorer Server software collected and displayed

The eHealth Live Exceptions module graphs activity levels and shows link detail for items you select in the tree view.

Visual UpTime's Access Utilization by Channel report identifies traffic levels by protocol for a WAN link.

for each data link connection identifier (DLCI), yet still used our network quite frugally to inform Visual UpTime of the network's current status. We found UpTime's calculations of round-trip delays very accurate. Those calculations excluded router serialization and insertion delay, and thus gave us a precise measurement of network delay for each permanent vir-

(via a browser interface) a wealth of statistics on the health of the WAN link and showed us who uses the most bandwidth and link utilization. Network protocol distribution reports showed the relative traffic levels of WAN protocols. The primary reports reveal line availability and SLA breaches (both summary and detailed versions). Other WAN link-related reports

show line statistics, DLCI traffic by bandwidth consumption, PVC by CIR load, DLCI performance and response times.

Network Instruments' Observer is more than just a protocol analyzer or packet decoder. It also can accumulate network activity statistics and display them in useful ways. When you put the vendor's hardware or software probes on remote network segments, Observer collects network activity statistics from those probes. Observer polls these probes every 5 seconds (by default), and you can increase this to every 2 seconds. Observer presents the latest, average and maximum overall bandwidth utilization statistics, maximum and average utilization by DLCI, top talkers and congestion metrics, which include notifications when congestion is occurring, even when bandwidth utilization is below the CIR. Observer also works with probes from other vendors, such as Netscout.

Concord's eHealth includes four modules — LiveHealth, Network Health, System Health and Application Health. Network Health monitors the performance and availability of WAN interfaces, routers, switches, frame relay circuits and remote access equipment. System Health monitors servers and selected (or all) clients to alert administrators to application performance problems, server crashes and disk space shortages. Application Health is a transaction-oriented collection of tools that help determine the cause of poor application response times. At a default of 5-minute intervals (or at a rate you can set), eHealth actively polls SNMP-aware devices to determine their status and displays the result in real time.

EHealth recognizes and understands more than 900 management information base (MIB) definitions. It uses these MIBs to determine device performance and availability. Initially, eHealth collects network activity and inventory data to build a normal network baseline. Thereafter, using a complex but configurable rules set, it

detects and highlights exceptional activity patterns, such as excessively high or low traffic through a router or switch port. The Network Health frame relay module efficiently and accurately collected network statistics from the DSU/CSUs in our WAN links. EHealth's many reports showed us WAN link data such as top talkers, packet discards, congestion, overall utilization and utilization by DLCI (average, minimum and maximum). We found that eHealth also understands and can monitor DSL connections.

Neon Software says you can use Cyber-Gauge to monitor Internet connections, but we found it also can keep an eye on private WAN links. CyberGauge is well suited to small networks and Apple Macintosh-based networks. Using SNMP, CyberGauge queries an IP address (a router, for example) at the other end of a WAN link as often as every second, and collects MIB II data. However, we found setting the interval rate to 10 or 15 seconds let CyberGauge gather useful statistics. It reported uptime and downtime in terms of the number of intervals the link was active, and showed uptime as a percentage. It showed total bytes inbound and total outbound, as well as utilization billing information expresses as average traffic levels for 5-minute periods. CyberGauge also displayed bandwidth utilization for the reporting period in percentage ranges.

Ease of use

Visual UpTime excels at helping administrators maintain WAN link details, locate link problems and track link activity. Clearly its designers carefully and thoughtfully focused on administrator productivity as they built UpTime's responsive and intuitive user interface to fit the workflow and individual tasks within a large network operations center. For example, the Network Configuration dialog is a central point for changing or adding networks, sites, access lines, ASEs and circuits. UpTime's ability to

TWO FAST?

Now the world's fastest network color printer
has a challenger right on its tail. The Xerox Phaser® 6250.
There's a new way to look at it.

Our newest color printer has speed in its DNA. And why not. Its sibling, the Xerox Phaser® 7300 large-format color printer, is still a first-place record holder in its class* with 30 ppm color, 37 black and white. But now it has a speedy new family addition, the Phaser 6250. It runs flat out at 26 ppm for both color and black and white letter prints with a first page out in just 12 seconds. Either one of these color speedsters eliminates the need for multiple printers. And both can get you to your finish lines in record time. To catch brilliant color in a flash, check out our full line of award-winning network printers. But don't blink, or you'll miss them go by.

Visit: xerox.com/office/race1960 Or call: 1-877-362-6567 ext. 1960

THE DOCUMENT COMPANY
XEROX

WAN tools

print a network configuration report that documented our work was icing on the cake. We never had to fumble around in the interface to locate the right window through which to update network details, troubleshoot a problem or produce (or schedule) the appropriate reports for our WAN links.

Adtran's N-Form user interface did not impress us. Its main administrative window distinguishes between user-oriented and server-oriented tasks. Selecting the Users tab in the administrative tool brought up windows in which we could create, change or delete users. The Servers tab similarly was a doorway into configuring N-Form's default SNMP settings, network utilization thresholds, e-mail identities, event history log and network event notifications. N-Form's Network Manager interface displays a hierarchical tree of network segments that identifies devices by address, type and status. An administrator can attach comments to each device's N-Form data to help make the tree's entries more meaningful. Network Manager can discover and display non-Adtran devices, but the tree's "type" column is relevant only for Adtran devices. The tree's "status" column, whose information is only as recent as the last SNMP polling sweep, only shows either "offline" or whether an e-mail notification is associated with a specific device. N-Form's Network Manager tree can be collapsed or expanded to help drill down to specific segments and devices.

In contrast, we found that Allot's WiseWAN WANXplorer has a well-designed tree-view interface that contains an intuitive and clearly presented display of network devices. We could move objects via drag-and-drop and sort columns of data by clicking on the column header. Right-clicking an entry displayed WANXplorer's easy-to-understand pop-up menus. Best of all, WANXplorer color codes currently set alarms to show a rising status (red) or a falling status (gray).

Observer uses a tree-view main window and multiple concurrently open child windows to show devices and events. Drilling down to get more data is simply a matter of double-clicking an item in the tree. Observer also displays a window containing a graphical view of network conversations. Alongside each conversation pair are statistics showing packet-to-packet delay times, retransmissions and lost packets. Clicking on a conversation pair drills down to a list of packets exchanged by the nodes. Each display of network activity is a child window that updates in real time, and you can have as many concurrent windows open as you wish.

While the other tools presented native Windows interfaces, Concord's eHealth server console used The SCO Group's Xvision PC X server. But growing accustomed to PC X takes only a short while. EHealth's expandable combination of tree-view window and associated detail

windows gave us quick access to network segment and device details, and current status. EHealth obviously is intended for large networks. For example, we found we could sort eHealth's display of network devices by IP address or class, which helped make working with populous segments much easier. Creating circuit-specific presentations of uptime and bandwidth utilization is easy with eHealth.

The CyberGauge interface defines simplicity. Entering device data involves choosing an interface type from a list (including frame relay, Ethernet, and serial) that CyberGauge detects on the router you point at. CyberGauge then lets you configure interface preferences and parameters, and how you want to display statistics. After you select one or more interfaces on a target router, clicking the Begin Monitoring button puts CyberGauge to work.

All six tools offer browser-based access to their reports and configuration data.

Discovering devices

Entering details about each WAN link isn't a task you need to do every day, fortunately, but each tool takes a different approach to the job. You explicitly tell Visual UpTime, WANXplorer, CyberGauge and Observer about each IP-addressed device at either end of a WAN link. In contrast, using IP address ranges you specify, N-Form and eHealth automatically discover WAN link devices. In our tests, eHealth's discovery process occurred daily, on a schedule we could set or, if we wished, interactively. During each sweep of the network, eHealth automatically discovered new or changed device information. EHealth eased the process of identifying network devices by letting us categorize network elements by class or IP address grouping. It then performed a discovery process to find those elements on the network.

Other considerations

All the tools we tested handled the various protocols we threw at them. We went a step further, however. TCP has an internal throttling mechanism that classically fails in the presence of other protocols. The mechanism senses overall TCP traffic levels to know when to throttle itself back, but the traffic-level detection ignores other protocols as it decides how many packets it can send in a "window" before expecting a response from its session partner. How well would the traffic-shaping tools work when we mixed high levels of TCP and other traffic on the network? Working at the application layer, the Adtran IQ 710 and Allot WiseWAN 401 sorted out the traffic jam quite nicely as they prioritized, for example, database transactions over e-mail.

We found that UpTime and eHealth scaled best. They both exhibited the capacity to handle a range of different network sizes, as well as a high degree of modular configurability.

All six products integrated well with HP's OpenView, emitting SNMP alerts (traps) that OpenView accepted and processed.

■ How we did it

Our test environment had three T-1 links, three frame relay links and a 384K bit/sec symmetric DSL link. The T-1 and frame relay links consisted of pairs of back-to-back DSU/CSUs and Cisco 3500 routers. The SDSL link consisted of Efficient Networks' SpeedStream 5851 DSL modems and a Nokia D50e DSL access multiplexer. The three frame relay links had committed information rates of 56K, 256K and 384K bit/sec. The seven links, singly and in combination, simulated increasingly complex WAN pathways among four 100-MHz Fast Ethernet network segments.

Our client platforms include Windows 98/ME/NT/2000, Red Hat Linux 6.2 and Macintosh System 8. The relational databases on the network were Oracle 8i, Sybase Adaptive Server 11.5 and Microsoft SQL Server 2000. Windows NT/2000 and NetWare 5.1 shared files, while Internet Information Server, Netscape and Apache software served up Web pages. The network's transport-layer protocols were TCP/IP, IPX/SPX, AppleTalk and SNA.

UpTime is a Windows-based monitoring tool the vendor delivers pre-installed on a Dell server. Observer is a Windows-based monitoring tool. N-Form is a modular, Windows-based, Java-enabled software tool for configuring and collecting data from IQ 710 DSU/CSUs. WiseWAN WANXplorer runs on Sun Solaris and Windows, and includes a bundled Sybase relational database. CyberGauge runs on Windows and, interestingly, Macintosh System 7 or later. EHealth, too, is multi-platform, available for HP-UX, Solaris, Windows NT and Windows 2000.

Except for Visual UpTime, which was pre-installed on its own server, we ran each vendor's software on a 4-way Compaq Proliant ML570 computer with 900-MHz Pentium III CPUs, 2G bytes of RAM, eight 18G-byte SCSI RAID drives and two NC3134 10/100 network adapters. The operating system platform was Windows 2000 Advanced Server. An Agilent Advisor protocol analyzer generated packets, and decoded and displayed network traffic.

We particularly wanted each WAN monitoring tool to alert us quickly and accurately to WAN link outages and problems. We looked for reports that helped us establish baselines, show available and unavailable devices, log device availability histories and identify trends. We tested for accurate, complete interpretation of network events and processing of SNMP management information bases emitted by network devices. Comprehensive traffic analysis, device discovery and multiple protocol support were important. We examined the products' reports to determine compliance with the terms of an service-level agreement. The ability to interface with a network management system, such as OpenView, was a plus. We factored in the ease with which we could administer the product.

We used various techniques to cause WAN link error events in the lab, including powering down specific devices, generating high traffic loads and introducing electrical interference, via a wired shunt, at the V.35 interface of the DSU/CSUs.

They're also easy to install. And kudos to Visual Networks for sending customers Visual UpTime pre-installed on a fast server.

Visual, Concord and Network Instruments all supply professionally written, clear and comprehensive documentation, as well as useful online help. Allot Communications' documentation is a 96-page manual that explains the WiseWAN hardware, but leaves the bulk of the software's description to the online help files. Adtran's documentation consists entirely of online help files, while the CyberGauge documentation is simply a 44-page booklet augmented by some online help files.

Conclusion

Migrating to Visual UpTime when you install new or replacement DSU/CSUs can help you create a WAN environment that's conducive to keeping your WAN links running smoothly. Troubleshooting the problem of the hour is much easier when you have Visual UpTime's level of up-to-the-second detail to help. To avoid future problems, Visual UpTime's reports are a godsend to capacity planners who need to make intelligent judgments about network

growth and changes.

For heterogeneous networks, eHealth is just what the doctor ordered. Its status indicators showing the condition of your network segments and devices — as well as its plethora of useful reports — make eHealth a necessity on large, diverse networks.

Nance, a software developer and consultant, is the author of Introduction to Networking, 4th Edition and Client/Server LAN Programming. *He can be reached at barryn@erols.com*

Fighting back
against telecom
surcharges

Understanding the distinction between regulatory charges
and the extras carriers pass along is a powerful weapon
when negotiating contracts.   ■ BY DAVID ROHDE AND STEPHEN SHEA

Running a corporate telecom shop often feels like playing the Whack-a-Mole machine at a carnival — you keep pounding down the moles but they keep popping back up. That's never more true than in the game of bill surcharges, in which every move by user groups or the government to rein in confusing bill add-ons soon is parried by an even more maddening and confusing line item from the carriers.

Don't be fooled. When a new line item pops up on your bill, it's all revenue to the carrier, no matter whether it's labeled as a tax, surcharge or pass-through.

From time to time the government tries to get the carriers to explain accurately who is responsible for the purported costs leading to each surcharge. But usually that just leads the carriers to break out some fees and combine others in a fashion that just barely complies with whatever the latest rule requires.

Your best bet is to understand why the carriers view the fees as an essential part of their revenue management — and then build these fees into the real cost of each of your network services, rather than thinking of them as government mandates, as the carriers proclaim.

That way you can bring your entire projected surcharge you spend into your carrier contract negotiations and apply the same principles of competitive leverage that you do with each service's regular rate elements such as tolls, ports, circuits and features. Tallying the real cost this way can help bring about overall contract concessions and credits, even if each surcharge is called "undiscountable" in the carrier's official service guide. But to add up these costs you need to know each carrier's surcharge platforms and what their plans are for them, because even some of the most obvious surcharges have new twists and turns.

Reform and results

Most users are familiar with the big surcharge for universal service, or universal

connectivity, but recent changes in the rules have altered the challenge in managing these expenses.

Last year the FCC declared it was fed up with what had become increasingly blatant markups of the universal service fee by long-distance carriers. For much of 2002, the FCC said it needed 7.28% of applicable carrier revenue to fund certain telephony and Internet subsidies. Yet by the fourth quarter business customers actually were being charged 8.3% to 9.6%. And what really peeved the FCC was that AT&T was charging residential customers 11%.

So the FCC decided that starting April 1, the carriers would have to start passing along the universal service charge with no markup or get rid of the line item and incorporate the cost into their service prices. None of the major carriers dropped the separate fee, but they did fall into line and now all business and residential customers receive the exact surcharge that the FCC mandates.

Except for three problems:

• The insatiable needs of the universal service program caused the FCC to increase the official fee — what it calls the quarterly contribution factor — to 9.1% for the second quarter of 2003 and then to 9.5% for the third quarter before bumping it back down to 9.2% for the fourth quarter. That alone is close to or above what business users previously paid some carriers after the now-illegal markup was applied.

• The FCC gave in to arguments especially by AT&T that the carriers incur "administrative costs" in collecting money for universal service. The government agency authorized carriers to separately recover these costs from users provided they didn't call this extra cost a regulatory fee or use other wording that seems to blame the government.

• The resulting administrative expense fee of 0.74% that AT&T initiated April 1 might have a clever name meant to avoid explicitly blaming a government rule. But this type of fee is hard for many users to distinguish from the existing AT&T federal regulatory fee, which AT&T subsequently raised on Aug. 1 from 0.35% to 0.52%. What's more, the just-increased federal regulatory fee now is buried in a combined line item with an unrelated surcharge regarding AT&T's property taxes.

MCI and Sprint customers face almost the same challenge. Both carriers charge nominal regulatory surcharges, but until recently had held out on adding an administrative fee relating to universal service fund (USF) collections. That also changed Aug. 1, when Sprint added a 0.03%

Surcharge summary

If you examine your corporate telecom bill closely, you'll find several regulatory-related fees.

Surcharge and purported need	AT&T	MCI	Sprint
Universal service fee: Fixed by federal decree.	9.20%	9.20%	9.20%
Administrative expense fee: Recently set by carriers for expense of carrying out USF dictate.	0.74%	None	0.03%
Federal regulatory fee: Historically set by carriers for general compliance with laws.	0.52%	0.43%	0.51%
Property tax fee: Set by carriers to offset their costs.	0.81%	1.40%	1.54%
Total	11.27%	11.03%	11.28%

Does not include federal line charges, federal excise tax or state-specific charges.

administrative fee. So far MCI is holding back on issuing such a fee but might feel less constrained after it emerges from Chapter 11 bankruptcy.

Decimal points add up

Don't let those decimal points fool you — the fees eat up dollars quickly. Business user advocacy groups such as the Ad Hoc Telecommunications Users Committee have warned the FCC that if it's going to let carriers charge administrative fees, the fees at least should be a percentage of the size of the programs they're supposedly administering. Right now they aren't.

Say a company spends $100,000 in a month on qualifying charges. Under the new rules it pays 9.2%, or $9,200, to the USF. So the extra administrative fee should be a token percentage of that $9,200, right? Instead, AT&T's administrative fee is 0.74% of the entire $100,000 — or another $740, an absurd number for the purported administrative expense of one line item on one customer's monthly bill.

The Ad Hoc Committee pointed the FCC to a long-standing practice with sales taxes, by which courts allow a percentage of administrative expense to be calculated only from the tax itself, not from the amount of the customer's purchase. Even

then, the administrative expense is deducted from the government's tax receipts, not added on as another retail cost.

The FCC is considering new ideas to discard the revenue-based system of universal service assessments in favor of one based on the number of phone lines, telephone numbers or connection bandwidth installed. But until that matter is taken up again, users can expect that administrative expenses will be an openended invitation to surcharge creep.

The carriers point to the government when assessing some other non-tax surcharges, often saying that government "mandates" cause the charges. They don't indicate that the carriers' own lobbyists promoted some of these "mandates" as replacement revenue for failing business.

For example, when was the last time you dropped any coins into a pay phone? Several years ago, the Bell companies and the country's many independent pay phone operators were alarmed at the number of times that people were striding up to pay phones and dialing toll-free access numbers for calling card schemes and alternatives to collect calling. So the pay phone operators lobbied to have the recipient of a toll-free call pay something for the cost of the call having been initiated at a

pay phone — and the concept was written into the Telecommunications Act of 1996. The FCC decided on a per-call surcharge of 26 cents, which has since raised by some carriers to as high as 47 cents.

The FCC thought the recipient would be the carrier on whose network the inbound toll-free call landed. But the called party ended up with the bill.

Some corporate users deal with this problem by blocking calls initiated at pay phones and eliminating calling cards in favor of corporate wireless deals. But this is no comfort if your customers expect to reach your call centers. And airlines and hotels have no practical option to block pay phone charges other than making it clear that this revenue — like all other carrier income — will be negotiated as part of the complete telecom deal.

Other surcharges are nothing more than cost recovery for software upgrades and other capital expenses that carriers don't want to shell out on their own, such as local number portability or telecom-related public safety initiatives.

Don't fall for carrier spin about surcharges as something they dislike, when they really are trying to put off reasonable network upgrades. Earlier this year wireless carriers stuffed bills with a "consumer

alert" pointing to a Web site claiming that "taxes, government fees and the cost of government mandates" can add as much as 20% to your bill. But those mandates were simply for wireless number portability and E-911 capability that wireless carriers should be happy to undertake if they're serious about selling their services as a substitute for wireline connectivity.

The best option for users is to trust the information from real user advocacy organizations such as the Ad Hoc Committee and consumer groups that are trying to limit surcharges to actually unavoidable carrier expenses that the government directly bills them for.

Until that day comes, ask for contract clauses that limit surcharges to those reasonably related to the carrier costs. And make sure your carrier negotiations start with the assumption that surcharges are part of the price proposals, rather than wait until the end of a contract negotiation to recognize these costs when it's too late to do anything about it.

Rohde is a senior analyst and Shea is a principal at TechCaliber Consulting, a telecom consulting firm. They can be reached at drohde@techcaliber.com and sshea@techcaliber.com.

Management Strategies

Planning for chaos

Learning to expect the unexpected helps bring order to project management.

■ BY CHUCK YOKE

The project appeared to have all the ingredients for success. Implementing a consolidated campus network would generate a 30% reduction in monthly telecom bills. Capital costs were less than $500,000, senior management endorsed the project and required resources, the work breakdown and milestones were documented and approved, and the critical path of key events necessary for the project to be completed on time had been identified.

Nothing should have gone wrong. Yet the project was completed over-budget and behind schedule. New hardware requirements for video connectivity and remote access were added at the last minute, resulting in cost overruns of nearly $200,000. Early snowstorms delayed the fiber installation, while equipment manufacturing and shipping backlogs caused implementation delays that eliminated three months of projected savings. By the time the network was finally live, it was a bittersweet success at best.

No matter how much we plan, unforeseen events appear to create havoc with a project's budget, scope or timeline. According to chaos theory, this is inevitable. For in our seemingly orderly world, one of the fundamental processes that pervade our universe is chaos. Philosopher Fredrich Nietzsche proclaimed "out of chaos comes order," and theorists such as James Gleick say chaos is the breeding ground of order.

Extrapolating the theoretical into the practical, it might be impossible to avoid chaos. So instead of trying to deny it, we should plan for it and even cultivate a controlled chaos in projects to reach the order we seek.

Five years ago, a typical project required this series of events: Identify a need, develop a solution, write and approve a business case, assemble a project team, develop a project plan, approve the budget and launch the project. A project was considered a success if it is was implemented on time and under budget.

The pace of business has quickened, and overall timelines have shortened drastically. The sequential events of yesterday need to be the parallel events of today. Business cases often need to be written while solutions are being developed. Project timelines need to be created before all the tasks and resources have been completely identified.

This new fast-paced paradigm is a breeding ground for chaos. To be successful, project managers need to be flexible, creative and able to respond to events quickly. Instead of assuming the solution will work, have an alternative plan in case it doesn't. Instead of assuming that key resources will be in place throughout the project, identify other people who can jump in and take over if necessary. Instead of assuming that the budget is final, know how to cut expenses by 10%.

My own philosophy is that a project manager should be like a sheepdog. Sheepdogs set up a boundary for the flock, allow a certain amount of freedom within that boundary and nip the heels of those who try to cross the boundary. Instead of trying to control the individual sheep, the sheepdog focuses on moving the entire flock to the stated goal.

I tried the command-and-control method of management early on in my career with little success. Instead of listening to my reports and asking for their input, I tried to tell them how to manage their systems, design their networks and implement security. After my arrogance led to the loss of some key employees, I realized that my staff wanted respect for their skills, acknowledgement of their ideas and a certain amount of ownership over their work.

I set boundaries, asked for input into key decisions, sought their expertise, gave them control over aspects of their job and allowed a certain amount of controlled chaos to reign in solutions development and implementation. As a result, both my success rate and team morale rose.

Recently, I was handed a VPN project that was already behind its initial time frame. In order to reap the anticipated cost savings, the solution needed to be implemented quickly. The short time frame necessitated creating the business case, engineering design and project plan in parallel.

Instead of trying to control every aspect of the project, I worked with the project manager to establish boundaries. We allowed a lot of chaotic movement within those boundaries and nipped the heels as needed to keep the team moving toward the goal.

We involved the project team in brainstorming various scenarios to develop risk-mitigation plans. We identified and documented the assumptions used, and worked to ensure project sponsors and senior management were aware of all timeline and budget risks.

During network planning sessions we cultivated controlled chaos and allowed a free exchange of ideas, sometimes playing devil's advocate to insure all options were reviewed. We facilitated intense technical discussions concerning vendor-provided vs. in-house networks, IP Security vs. generic routing encapsulation VPNs, and pre-shared keys vs. certificate authorities. When discussions wandered off the track, became personal or overran the project scope, we would nip heels as needed to bring everyone back into focus.

Even after a decision had been made on hardware, our timeline contingencies let us accommodate additional chaos and incorporate newly available equipment from a different vendor that brought added value, but required a redesign of ordering, shipping and configuration processes.

From outside, the project often appeared to be in total disarray. There were times when the proposed design changed on a weekly — and sometimes daily — basis. Many times the project manager or I had to meet with the project sponsors to assure them everything was OK.

We let chaos reign, but it was controlled chaos. Like sheepdogs, we constantly circled the team making sure that timelines were met, solutions were in scope and equipment was in budget. Ultimately, out of the chaos came order.

By both planning for chaos and cultivating controlled chaos, within 90 days of final budgetary approval we implemented a solution that not only met all customer requirements, but also provided a scalable platform that would accommodate network growth. Nietzsche would have been proud.

Yoke is a business solutions engineer for a corporate network in Denver. He can be reached at ckyoke@yahoo.com.

More online! Chuck Yoke explains his sheepdog management philosophy for moving your flock toward the desired goal. **www.nwfusion.com DocFinder: 7921**

ALLEN CRAWFORD PLANKTON ART CO

Advertising Supplement

IT Careers: The Next Generation Data Center

There was a time when data warehousing and data centers were primarily a hardware issue or deep research embedded within data management. Not so with the Next Generation Data Center, where data mining and statistics come out of the research cave into full blown operational imperatives. You've seen the results – a list of books that you just might find interesting, based on your past orders.

With more than one million people working in the data management field, the new role for data mining and business intelligence indicates a major shift of web development, software, data mining (statistics) and text mining into the job category. This segment of the information technology industry is moving up in the food chain; the analytics category is considered one of the top value plays with over half of all implementations achieving full payback within two years or less. And, expenditures are expected to grow by better than 33% over the next four years.

Among the kingpins of mining and putting to work data to achieve bottom line results is Amazon.com. Ken Collins, director of data warehouse for Amazon.com, says the most critical issues are data quality, scaling and "deployment of mining insight throughout our infrastructure. I would expect that companies that can make deployment of sophisticated mining models and algorithms across an n-tier open systems stack would do very well."

Wayne Thompson, product manager for data mining technology at SAS Institute in Cary, N.C., says that in the past 24 months, data management has shifted from a tool kit product to a solution of best practices resulting in new business capabilities "such as credit scoring, bioinformatics, marketing automation and money or fraud identification. It's now a solution within a larger enterprise," says Thompson, "which allows more focus and a more consumable product." SAS, long known for its analytical focus, holds a 38% share of the data mining market.

In the past, data mining resulted in quantitative points – age, gender, zip code, how many cars you own. The changes Thompson sees are quantifiable and predict future behavior. Another advancement is the use of statistics with text. "Whether in business or research, much of the information we receive is in the form of text. At SAS we've come up with a way to combine textual with typical numeric data to come up with better predictive models," Thompson says.

The industry is also seeing data management move from research into a collaborative work environment. Working with business leaders and the IT professionals who deploy data systems within larger enterprise or customer relationship management systems, data management professionals are developing models based on business trends.

As these shifts occur, the shift in skills required is equally broad. The collaborative environment requires written and verbal communication abilities, as well as business understanding. The requirement for solid statistical skills (as evidenced by master's or doctorate degrees in mathematics, statistics, computer science or even data mining) remains. Standard languages – PMML (predictive modeling markup language), which is an XML representation – is the base, but experience in C, C++ and JAVA also are important. "We also need people who can sit with a business manager, understand the business processes and develop models," Thompson says.

As with other IT job categories, the opportunities exist in IT companies such as SAS. They also are increasingly available in non-IT firms who are customizing and fully exploiting the data mining tools to analyze and predict. Thompson also recommends direct exposure to data management/mining tool kits, such as SAS Enterprise Miner or Text Miner, and attending conferences such as M2003, to be held Oct. 13-14 in Las Vegas. "It's an excellent forum that focuses less on statistical skills than on how tools are being used to affect business," he says.

Amazon.com's Collins adds, "If ever there was an opportunity for IT to be proactive for business, this is it. Professionals must understand data mining and statistics, their customer or company's business and help map the two together. It is also essential that we be relentless in the pursuit of data quality and metadata for the high business value information we steward. And of course, nothing replaces solid data warehouse modeling, design and optimization skills."

Next Generation Data Center
Mines text as well as data points
Combines business modeling with analytics
Predicts rather than reports

Next Generation IT Skills
Advanced analytical skills (statistics)
Written/verbal communication
Business models and best practices by industry sector
Software development
Enterprise-wide and web-enabled deployment

For more information about IT Careers advertising, please contact:
Nancy Percival
Vice President, Recruitment Advertising
800.762.2977
500 Old Connecticut Path
Framingham, MA 01701
Produced by Carole R. Hedden

Looking For A New Career?

The new

itcareers.com

and

CareersJournal.com

combined

jobs database

can help you

find one.

Check us out!

www.itcareers.com

System Software Engineer - Work as part of software development team to write software for automatic ultrasonic imaging systems, including motion control, instrument setup, data acquisition, operator interface, image display and processing, and data analysis. Focus on instrument control and data acquisition software and graphical user interface design. Interface with customers, in-house sales staff, hardware development engineers and manufacturing engineers to specify, customize and implement software modules. Provide accurate estimates of required time to complete software tasks. Use Microsoft SourceSafe for version control. Provide on-site support for installation and testing of system software as required. Interface with customers as required to assist with support and service. Program in C/C++ and Visual C++ under Windows NT/98. Use object oriented design, motion control software and GUI knowledge. Requirements include a Bachelor's degree or equivalent in an Engineering discipline, Physics or closely related technical field and three years of work experience in the job offered or related field of system and control software engineering. Applicants must have unrestricted authorization to work in the United States. Salary $74,838/year. 40 hours/wk. Respond with two copies of resume to Case #200202405, Labor Exchange Office, 19 Staniford St., 1st Fl., Boston, MA 02114

Software Engineers: Software life cycle development in Microsoft Development tools suite including C#.Net, Oracle RDBMS tools suite, SQL Server, DB2, COBOL and related tech. to develop array of apps. including secure Client Server/Web based HR, Communication and Government apps. Complete job description or to apply contact HR, 3761 Venture Dr., Bldg. 100, Suite 240, Duluth, GA 30096.

MECHANICAL ENGINEER

Require BS in mechanical engineering and 3 yr. exp. in position offered. Must have taken courses in thermodynamics, hydromechanics, and heat transfer theories and research exp. in energy field. Resume to Arbin Instruments, 762 Peach Creek cut Off Rd., College Station, TX 77845.

Java Programmer Analyst. Provide technical computer support and user assistance in developing, operating, and maintaining all web application programs for all divisions of Morley Companies, Inc; maintain existing web applications; and prepare program documentation. Must have Bachelor's in Computer Science or related, and knowledge of AS400 platform, Web Sphere Application Server, and HTTP server for IBM, EJB, Java, JSP, & Javascript.. Send resume: Morley Companies, Inc., Attn: Richard Mott, One Morley Plaza, Saginaw, MI 48603.

SOFTWARE ENGNR: Analyze, dsgn & dvlp web-based applns & systems. Specific duties include: (i) dsgn, dvlp & config-ure websites & systems per client specs using VB, ASP, MTS, Visual Source Safe, Java, Visual Interdev, DCOM, COM+, Active X, ADO and IIS; (ii) act as Program Lead for client's projects; (iii) perform system analysis, coding & testing; (iv) dvlp components to integrate Call Center OA tools w/Kronos; (v) maintain & perform sys. optimization & compliance w/system reqmnts of health care industry; (vi) perform intensive data analysis & dsgn; & (vii) dvlp identity & account mgnt systems based on client's reqmnt & spec. Bach. degree in Comp. Sci. or Electrical Engmg + 3 yrs exp. in position offered or as a Prog. Analyst or Software Dvlpr/Consultant reqd. Must have: (1) 2 yrs of exp. w/VB, Java, J2EE, JavaScript, VB Script, ASP, COM, COM+, Active X, ADO, XML, SOL Server, SOAP, Seagate Crystal Reports, KRONOS, Visual Interdev & Visual Source Safe as well as working w/iBaaN CRM components (incl. Baan EIA and BOIS as well as iBaan CRM); (2) 2 yrs of health care industry exp., including exp. w/compliance stndrds (HL7, CCOW, HIPAA); (3) ability to adhere to industry-wide software dvlpmnt stndrds, incl. SEI-CMM, SEI-PSP; & (4) high mobility preferred. 40 hrs/wk, 8 am - 5 pm, $64,240/yr. Qualified applicants please submit resume to: McKeesport CareerLink, Attn: ES Manager, 345 Fifth Avenue, McKeesport, PA 15132-2600. Please refer to Job Order No. WEB 360618.

ObjectWin is looking for Programmer/System Analyst or other IT professionals. Applicants must have BS or equivalent. Skills in ASP.Net, B2B, VB, Java, HTML, VB.Net, XSL, CSS, MS CMS, SSL & MS certified preferred. Competitive wage. Apply at: skarande@objectwin.com. EOE.

Techgene Solutions has openings for System Analysts or Software Engineers. Candidates must have BS with experience in Cobol, JCL, Oracle, SOL, etc. Travel may be required for some positions. We are small but stable. Competitive salary. Please apply at bapujik@yahoo.com. EOE

OmniPros, a worldwide provider of software solutions seeks motivated IT professionals specializing in: Java, J2EE, Oracle, WebLogic, Webmethod, Vitria, Tibco, Portals, Oracle CRM (Technical, 11i) Oracle Mfg (Technical/Functional, 11i), Oracle Finance (Technical, 11i), Business Development/ Technical Operations Management Positions are located in Chicago, Il. Please e-mail resume to careers@omnipros.com, fax resume to 408-944-0719, or mail resume to: OmniPros Ltd. 99 W. Tasman Drive, Ste 205 San Jose, CA 95134

Systems Administrator sought by computer s/w development firm in Jacksonville, FL. Must have Bach in Comp Sci., Engg or equiv and two yr relevant exp in designing, developing and implementing LAN, WAN, VPN and telephone networks; Windows/Linux-based DNS, DHCP, WINS Servers/Domain controllers; SOL Server7.0/2000; and Oracle 8i/9i for OLAP & OLTP databases. Respond to: HR Dept., Intelligenxia, Inc., 4905 Belfort Rd., Ste 110, Jacksonville, FL 32256.

Senior Software Engineer-Platform Services: Lead and participate in specification, design, development and support of company products including the overall architecture, component interfaces and communication schemes, client and server-side programs written in Java and C++. Develop Oracle, SOL Server and LDAP database schemas. Assist with development of project plans and schedules. Follow rigorous software engineering standards including developing product requirements, functional and design specifications and adhering to coding standards. Lead efforts to identify and resolve any product performance issues. Mentor junior engineers. Requirements include a Master's degree or equivalent in Computer Science, an Engineering discipline, Mathematics or related field and three years of work experience in the job offered or related field of software engineering using C++, or a Bachelor's degree or equivalent in Computer Science, an Engineering discipline or related field and five years of progressively responsible experience in the job offered or related field of software engineering using C++. Applicants must have unrestricted authorization to work in the United States. Salary $78,000/year. 40 hours/wk. Respond with two copies of resume to Case #200203134, Labor Exchange Office, 19 Staniford St., 1st Fl., Boston, MA 02114

Software Engineer - Provide software engineering consulting services to clients for Storage Area Network (SAN) systems that may include systems manufacture by Compaq, EMC and/or Hitachi, or others. Design and implement management and monitoring SNMP software for the SANs and other IP devices. Configure and test SAN devices including, but not limited to: Brocade Switches, Foundry 68000 routers, Crossroad routers, Cisco and Tape Libraries. Will use software languages including Java, CORBA, HTML, Java servlets, Visual Basic, Pascal, JSP, SML, JMS and networking protocols such as SNMP, TCP/IP, HTTP, HP Overview, etc. Work includes SAN/LAN programming for high scalability. Masters degree in CS and two years of experience or Bachelor's degree and five years of experience on software work including SANS systems. Experience must include: Design, develop, implement, test SANs systems; work with routers and switches such as Cisco routers/switches, Foundry 68000 switches; work on SAN/LAN solutions; Must be willing to be assigned to unanticipated client sites throughout the United States. Hours: M-F, 8:00 a.m. to 5:00 p.m.; 40 hrs/wk. Salary: $85,384/year. Send duplicate resumes to: Job Order # 2003-336, P.O. Box 989, Concord, NH 03302-0989.

Seeking qualified applicants for the following positions in Memphis/Collierville, TN: Senior Technical Analyst, Research, evaluate, implement and coordinate changes to computer systems/applications. Requirements: Bachelor's degree or equivalent in computer science, math, engineering or related field plus 5 years of experience in systems/applications development, including programming. Experience with IMS DB/DC, Cobol II and TSO also required. *Master's degree in appropriate field will off-set 2 years of general experience. Submit resumes to Sibi George, FedEx Corporate Services, 1900 Summit Tower Blvd., Suite 1400, Orlando, FL 32810. EOE M/F/D/V.

Regional Sales Manager   West Coast

IT Careers has an exciting opportunity for a highly motivated and seasoned sales professional to join our team. This Regional Manager will be responsible for selling integrated recruitment advertising packages consisting of Print (within our network of publications), Online and Recruitment Event Show Booths, as well as other products as they become established. Emphasis will be on generating revenue by developing new accounts, but this individual will be responsible for growing existing accounts and servicing recruitment ad agencies ensuring that IT Careers is top of mind. The candidate will be the part of a sales team working to set the territory strategy, therefore collaborative selling skills are critical. In addition, this person must have a solid understanding of the IT recruitment market, be an effective communicator and negotiator and have a proven record in sales. Minimum of 5+ years of outside sales experience required, preferably in recruitment advertising sales. This position will be based in the candidate's home office on the west coast. Travel is required.

If interested, please email a resume to jcjobs@idg.com or fax to (508)935-4600. Please include code ITC922 in the subject line.

ITcareers

Software Developer - Design code and test financial models and their integration into a trading environment. Analyze, review, and rewrite program logic to increase operating efficiency to adapt programs to new requirements and migrate them as new standards arise. Design and test trading applications and write and rewrite program logic. Use experience in C++, Visual Basic, Excel and database administration. Will utilize basic knowledge of finance and economics. Requirements includes a Bachelor's degree or equivalent in Computer Science, an Engineering discipline or related field and two years of pre- or post-degree work experience in the job offered or related field of software development. Experience must include use of C++, Visual Basic and database administration. Applicants must have unrestricted authorization to work in the United States. Salary $75,000/year. 40 hours/wk. Respond with two copies of resume to Case #200203201, Labor Exchange Office, 19 Staniford St., 1st Fl., Boston, MA 02114

Software Engineers & Programmers: Analyze, design, test and implement specialize software applications for e-commerce, Web, Client Server technologies, Legacy systems and distributed apps. in Weblogic, Corba, Apache, Mainframe, ASP, J2EE, Siebel, PB and related technologies utilizing appropriate RDBMS including Oracle and DB2. HR, Instcomp, Inc., 906 Lacey Ave., Suite # 206, Lisle, IL 60536. EOE.

SW Dev Engineer (Omaha, NE) - Partner in a team environment to design, develop, code, test & debug new SW or significant enhancements to existing SW using Oracle databases, PL/SQL, WebLogic Server, Java, JDBC, & the J2EE/JDK environment. Write Java servlets using EJB, JSP, & JavaScript. Perform maintenance on existing SW using Oracle, PL/SOL, Java, & HTML in Windows NT, UNIX, & Linux environment. Apply principles, theories & concepts & use methodologies, tools, doc processes & test procedures to complete projects. BS Comp Sci, Eng, or related + 2 yrs related exp + working knowledge of: Oracle; PL/SOL, HTML, WebLogic Server, Java, J2EE, JDBC, EJB, JSP, Java Script, Java Dev Kit (JDK), Java Servlets; & Win NT/UNIX/Linux. $72,580/yr. M-F, 8-5. Must have proof of legal authority to work permanently in U.S. No calls. Send cover ltr & resume to: Madhavi Bhadbbade, Coordinator, NE DOL, P.O. Box 94600, Lincoln, NE 68509-4600. Ref. Job#TREL5-RPS5A-01300.

Systems Analyst. B.S. in Comp. Sci. or equiv. + 3 yrs. rel. exp. Exp. to include C/C++, UML, Python, XP, JavaScript, SQL, STL, MacApp, Mac OS X (carbon), and .Net. Send resumes to Robert Long, MetaCommunications, Inc., 1210 S. Gilbert St., Iowa City, Iowa 52240-4506.

NWG31006E/W/MW 3

Sales Offices

Carol Lasker, Associate Publisher/Vice President
Jane Weissman, Sales Operations Coordinator
Internet: clasker, jweissman@nww.com
(508) 460-3333/FAX: (508) 460-1237

New York/New Jersey
Tom Davis, Associate Publisher, Eastern Region
Elisa Della Rocce, Regional Sales Manager
Agata Joseph, Sales Associate
Internet: tdavis, elisas, ajoseph@nww.com
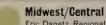(201) 634-2300/FAX: (201) 634-9286

Northeast
Donna Pomponi, Regional Sales Manager
Caitlin Horgan, Sales Assistant
Internet: dpomponi, chorgan@nww.com
(508) 460-3333/FAX: (508) 460-1237
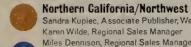
Mid-Atlantic
Jacqui DiBianca, Regional Sales Manager
Marta Hagan, Sales Assistant
Internet: jdibian, mhagan@nww.com
(610) 971-1530/FAX: (610) 975-0837

Midwest/Central
Eric Danetz, Regional Sales Manager
Agata Joseph, Sales Associate
Internet: edanetz, ajoseph@nww.com
(201) 634-2314/FAX: (201) 712-9786

Northern California/Northwest
Sandra Kupiec, Associate Publisher, Western Region
Karen Wilde, Regional Sales Manager
Miles Dennison, Regional Sales Manager
Courtney Coughlin, Regional Sales Manager
Maricar Lagura, Office Manager/Sales Assistant
Teri Lowe, Sales Assistant
Internet: skupiec, kwilde, mdennison, ccoughlin, mlagura, tlowe@nww.com
(510) 768-2800/FAX: (510) 768-2801

Southwest/Rockies
Becky Bogart Randell, Senior District Manager
Angela Norton, Sales Assistant
Internet: brandell, anorton@nww.com
(949) 250-3006/FAX: (949) 833-2857

Southeast
Don Seay, Regional Sales Manager
Caitlin Horgan, Sales Assistant
Internet: dseay, chorgan@nww.com
(404) 845-2886/FAX: (404) 250-1646

Customer Access Group
Tom Davis, Assoc. Publisher Eastern Region/General Manager, Customer Access Group
Shaun Budka, Director, Customer Access Group
Kate Zinn, Sales Manager, Eastern Region
Caitlin Horgan, Sales Assistant
Internet: tdavis, sbudka, kzinn, chorgan@nww.com
(508) 460-3333/FAX: (508) 460-1237

Fusion
Alonna Doucette, Vice President Online Development
James Kalbach, Director, Online Services
Stephanie Gutierrez, Online Account Manager
Debbie Lovell, Online Account Manager
Kristin Douglas, Online Operations Manager
Lisa Thompson, Online Ad Traffic Coordinator
Internet: adoucette, jkalbach, sgutierrez, dlovell, kdouglas, lthompson@nww.com
(610) 341-6025/FAX: (610) 971-0557

MARKETPLACE
Response Card Decks/MarketPlace
Jayson Cooper, Director of Marketplace Advertising
Enku Gubaie, Senior Account Manager
Amie Gaston, Account Manager
Chris Gibney, Sales Operations Coordinator
Internet: jcooper, egubaie, agaston, cgibney@nww.com
(508) 460-3333/FAX: (508) 460-1192

IT CAREERS
Vice President, Nancy Percival, Midwest/West Regional Manager, Laura Wilkinson, Midwest/West Account Executive, Mark Dawson, Eastern Regional Manager, Jay Savell, Eastern Account Executive, Andrew Haney, Sales/Marketing Associate, Joanna Schumann
(800) 762-2977/FAX: (508) 875-6310

Network World, Inc.
118 Turnpike Road, Southborough, MA 01772
Phone: (508) 460-3333

TO SEND E-MAIL TO NWW STAFF
firstname_lastname@nww.com

Evilee Thibeault, CEO/Publisher
John Gallant, President/Editorial Director
Eleni Brisbois, Administrative Planning Manager
FINANCE
Mary Fanning, Vice President Finance
Paul Mercer, Finance Manager
Betty Amaro-White, Event Finance Manager
HUMAN RESOURCES
Elizabeth Price, Director of Human Resources
Eric Cormier, Human Resources Representative
MARKETING
TerryAnn Croci, Senior Director of Marketing
Barbara Sullivan, Senior Research Analyst
Judy Schultz, Senior Marketing Designer
Cindy Panzera, Graphic Specialist
GLOBAL PRODUCT SUPPORT CENTER
Nancy Sarlan-Parquette, Sr. Product Marketing Manager
PRODUCTION SERVICES
Greg Morgan, Senior Director, Production Services
Karen Wallace, Senior Director, Advertising Operations
Mike Guerin, Senior Production Specialist
Jami Thompson, Production Coordinator
Veronica Trotto, Advertising Coordinator
Maro Eremyan, Advertising Coordinator
CIRCULATION
Richard Priante, Senior Director of Circulation
Bobbie Cruse, Subscriptions Manager
Mary McIntire, Circulation Marketing Manager
RESEARCH
Ann MacKay, Research Director
DISTRIBUTION
Bob Wescott, Distribution Manager/(508) 879-0700
IDG LIST RENTAL SERVICES
Paul Capone, Account Executive
P.O. Box 9151, Framingham, MA 01701-9151
(800) 343-6474/(508) 370-0825, FAX:(508) 370-0020
SEMINARS, EVENTS AND IDC EXECUTIVE FORUMS
Robin Azar, Vice President of Events
Michele Zarella, Director of Operations
Dale Fisher, Event Planner
Tim DeMeo, Senior Operations Specialist
Kristen Kennedy, Event Coordinator
Sandra Gittlen, Events Editor
Neal Silverman, Senior Director of Event Sales
Andrea D'Amato, Sales Director/Strategic Partnerships
Kristin Ballou-Cianci, Senior Event Sales Manager
Sandy Weill, Event Sales Manager
Maureen Riley, Event Sales Manager
Judy Tyler, Sales Operations Specialist
Mark Hollister, Senior Director of Event Marketing
Debra Becker, Dir., Marketing & Audience Development
Sara Evangelous, Marketing Manager
Timothy Johnson, Marketing Specialist
ONLINE SERVICES
Alonna Doucette, Vice President, Online Development
Hillary Freeley, Director, Online Operations
Deborah Vozikis, Design Manager Online
Adam Gaffin, Executive Editor, Online
Melissa Shaw, Managing Editor, Online
Jason Meserve, Multimedia Editor
Sheryl Hodge, Online Copy Chief
CUSTOMER ACCESS GROUP
W. Michael Draper, V.P. of Technology & Business Operations
Sharon Stearns, Mgr. Customer Access Operations
Cara Peters, Project Support Specialist
INFORMATION SYSTEMS/BUSINESS SERVICES
W. Michael Draper, V.P. of Technology & Business Operations
Anne Nickinello, Senior Systems Analyst
Puneet Narang, Manager of Database Technologies
Tom Kroon, Senior Software Engineer/Architect
William Zhang, Senior Software Engineer
Rocco Bortone, Senior Network Manager
Peter Hebenstreit, Senior Network/Telecom Engineer
Kevin O'Keefe, Systems Support Manager
Brian Wood, Senior Systems Support Specialist
Frank Coelho, Senior Manager, Business Services
Mark Anderson, Business Services Supervisor
Linda Cavanagh, Business Services Administrator

IDG
Patrick J. McGovern, Chairman of the Board
Pat Kenealy, CEO

Network World is a publication of IDG, the world's largest publisher of computer-related information and the leading global provider of information services on information technology. IDG publishes over 275 computer publications in 75 countries. Ninety million people read one or more IDG publications each month. Network World contributes to the IDG News Service offering the latest on domestic and international computer news.

NetworkWorld

Editorial Index

Advertiser Index

These indexes are provided as a reader service. Although every effort has been made to make them as complete as possible, the publisher does not assume liability for errors or omissions.

*Indicates Regional Demographic

IDS

growing, though they've been slow to catch on with buyers.

The Gartner report prompted such an intense argument among IT officials at the Department of Defense about buying IDS that the Office of the Secretary of Defense organized a meeting at the Pentagon in July. IT representatives and procurement officials from the Army, Navy, Air Force, Federal Aviation Administration, and departments of Energy, Justice and Homeland Security were also in attendance. Also included were a handful of IDS vendors and analysts.

Stiennon had no idea he'd be facing such a crowd.

"I didn't know the industry vendors would also be there," he says. "As I was walking down the hall to the room, they let me know."

According to meeting participants, Arbor Networks, Internet Security Systems (ISS), NFR Security, NetForensics and Sourcefire had been invited to represent the IDS point of view. In addition, two independent analysts, Greg Shipley, CTO at consultancy Neohapsis, and Peter Kuper, industry analyst at SG Cowen, were part of the roundtable discussion.

After Stiennon presented his "IDS is dead" arguments, he quickly came under attack by government personnel who had bought IDSs and were having to explain

■ **Read next week's issue for results of our "in the wild" test of four IDS products.**

their purchases to procurement officials, as well as industry vendors exasperated that Stiennon was making such a sweeping condemnation.

"People were saying 'Gartner makes statements about tracking hype, but who tracks Gartner?' Another said Gartner had an agenda to grab press,'" Shipley says of the meeting.

But Gartner's criticism struck a nerve with IT staff struggling to make IDS work and still dealing with worms and other threats, especially with internal software requiring patching."The Pentagon personnel were saying, 'We spend all this money on this security software and we still have problems,'" Shipley says.

Stiennon "was a little ganged up on," Kuper says, adding that he found Gartner's report on IDS to be "alarmist," "irresponsible" and based on outdated information about IDS technology, which he says is improving.

Kuper notes that the Gartner report might be having a freezing effect on IDS spending as IT departments are pressed harder to defend buying such products. But he also doubts customers would rush to buy firewall-based IPS offerings if they are already worried about false alerts with IDS.

As for the debate, little has been resolved.

"The Gartner guys aren't wrong in the issues they identified," says Marty Roesch, president of Sourcefire, and creator of the open source IDS software Snort. Roesch, who attended the meeting at the Pentagon, acknowledges that false alerts are a problem the industry needs to address. But, he adds, Gartner is "wrong in their conclusions. To

IDS or IPS?

Intrusion-detection and intrusion-prevention systems each have pros and cons.

	Pros	Cons
Intrusion-detection systems	Identify attacks, penetrations; useful for auditing, forensics.	Won't stop attacks; can raise false positives.
Intrusion-prevention systems	Can block network attacks; can be used in passive IDS mode.	Legitimate traffic can be blocked accidentally; as in-line devices, are potential points of failure.

recommend you don't need IDS anymore is ludicrous."

Shipley also defends IDS — to a point.

"Before you say they've failed, ask what did you intend them to do?" he says. As passive-monitoring systems, IDSs — fostered two decades ago through Defense Department research money — are primarily for auditing purposes, Shipley says.

In contrast, a firewall — the preferred Gartner approach — is "an enforcement device." He says it's not a clear-cut case that it makes sense to drop IDS for a firewall-like IPS that blocks traffic. There's a role for both.

Stiennon says his report erred in saying IDS products don't work over 600M bit/sec, as such systems now are reaching 750M bit/sec and higher.

While the Pentagon declined to comment on the IDS showdown or how future IDS and IPS purchasing might go, the meeting in July ended with no clear winner, according to several attendees.

"Defense Department people ended by summing up saying there's no clear decision today, but they don't like these false positives and 24-7 monitoring with IDS," Stiennon says. He adds that IDS vendors — many of which are adding IPS equipment to their lineups — now tell him that many government agencies in their RFPs are requiring in-line blocking at least as an option.

Customers following debate

Customers in the private sector are monitoring the debate with great interest.

"We see relatively few false positives," says Roger Safian, information security coordinator at Northwestern University in Evanston, Ill., which is using Lancope's StealthWatch IDS appliance to monitor network traffic in and out of the university's network. Asked if he would consider blocking attack traffic with an IPS,

he said, "I'm worried IPS will block legitimate traffic as well."

One concern with using an IPS is that a knowledgeable attacker could "figure out how to turn off your network" by tricking a device into blocking everything, says John McEachen, associate professor of electrical engineering at the government's Naval Postgraduate School in Monterey, Calif., which uses StealthWatch IDS.

He also notes that the U.S Pacific Command, based in Hawaii, is using a version of StealthWatch with a graphics-network-display addition called Therminator.

McEachen says military training calls for reliance on an "active watch standard" based on "human cognition as to what to do next." He says that means "human operators make decisions" when it comes to network attacks. Shifting to an IPS-based perspective in the military would entail change, but he says IPS could be seen as complementary to IDS.

"I'm a Gartner customer," says Andrew Conte, director of IT and chief information security officer at Home Box Office in New York. While the "IDS is dead" report has been food for thought, he says he's not ready to throw out his IDS for an IPS, which he sees as somewhat "immature" in terms of technology and market scope. "And you may be blocking valid traffic using IPS," he adds.

Paul Samadani, director of corporate technology services at Pentair, a tool maker in St. Paul, Minn., that uses the Sourcefire

IDS, also is leery of IPS.

"I just don't think we're at the place [where] we can do this well technically," he says, adding he questions Gartner's advice on this score.

However, many vendors that have their roots in passive IDS monitoring, including Sourcefire, also are developing products that can handle active blocking. Although it would mean designing an in-line device that analyzes by mirroring traffic and will stop attack traffic.

IDS vendor Intrusion last week introduced its first IPS sensor, SecureNet Sensor 5.0. And next week ISS will take the wraps off its Proventia line of multi-use IPS appliances at an event at Fox Electronics in San Jose. Gartner's Stiennon is expected to be there, along with Howard Schmidt, chief security officer of eBay and former White House security adviser.

A vote of confidence on IDS comes from computer forensics software maker Guidance Software, which this month is adding the ability in its Enterprise Edition 4.16 to capture data instantly, based on an IDS alert from Internet Security Systems and Enterasys products."We have faith in IDS, with fine-tuning," says Jon Blair, Guidance's senior

> **❝ IDS is like having a camera on the side of the highway. IPS is like a toll booth stopping the traffic. ❞**
>
> **Richard Kagan**
> Vice president of marketing, Fortinet

director of product development.

"IDS has gotten a bad rap," says Richard Kagan, vice president of marketing at Fortinet, which sells IDS and IPS products. "Gartner's entire argument is malformed. IDS is like having a camera on the side of the highway. IPS is like a toll booth stopping the traffic. They're entirely separate things." ■

More online!

Keep up to date on the latest trends in the security realm with our research page.

IBM to roll out next phase of ID management

■ BY JOHN FONTANA

IBM this week is scheduled to unveil upgrades to its identity management platform that more tightly integrate its suite of products and provide users with more automated controls of business workflow and applications.

The company has focused on IBM Tivoli Access Manager 5.1 for access control, Identity Manager 4.5 for provisioning services and Directory Integrator 5.2 to pull together repositories of user data.

IBM is tightening the integration between those products and others in its suite, including its Privacy Manager and Directory Server, to create an identity management platform that provides authentication, access management, user management and directory services. The goal is to eventually extend the entire package to support federated identity management among organizations using Web services standards.

The identity management initiative is part of IBM's $10 billion On-Demand strategy for enterprise computing. IBM is competing with rivals such as Microsoft, Novell and Sun that are working on their own comprehensive platforms.

In June, Sun upgraded its Sun One identity suite to include integration with Microsoft's Active Directory. In July, Novell unveiled its Identity Automation Framework, which incorporates its nSure product line. And later this month, Microsoft is scheduled to further flesh out its identity management platform, built around Active Directory and Identity Integration Server.

ID management is catching on

The management of user identities is becoming a hot project for end users because of the security and cost savings it promises.

"We said we would save $1 million-plus this year with identity management, and we have already met our goal," says Ronda Kiser, senior manager for enterprise automation for Whirlpool in Benton Harbor, Mich. "We had 11,000 calls to our help desk last year, and 60% to 70% were for password resets and account modifications."

Whirlpool rolled out IBM Tivoli Identity Manager early this year and recently upgraded to Version 4.5 to support self-service capabilities on its network, including self-registration. The company also has deployed IBM Tivoli Access Manager to support single sign-on for its users.

The next task is to further automate provisioning of user accounts with Identity Manager 4.5. Kiser says the provisioning improvements will save another $1 million for the company.

"As we centralize all our identity needs we get simplification, flexibility and we save money," she says. "But the real reason we are doing this is to improve security."

IBM has the same thought in mind with its upgrades to Access Manager 5.1, which it says will be available next month. The new Dynamic Rules Engine will let users pull additional user attribute information, such as age or credit rating, from a number of sources and apply it to authorization policies to tighten access controls.
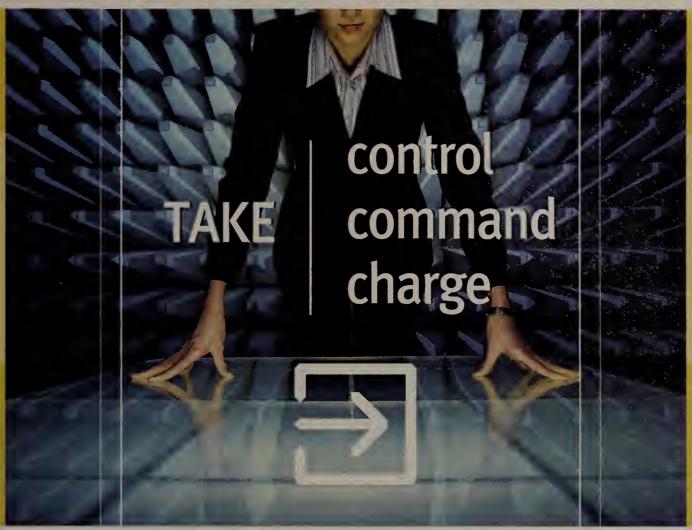
IBM also has added Dynamic Group Support, which lets companies assign access-control rights based on organization, job or partner status. The features also have been added to Privacy Manager 1.2, to support real-time checks on compliance with access policies.

With Identity Manager 4.5, IBM has opened the workflow engine so that it can be integrated with other workflow engines. Now identity management can be tied into larger business processes, such as setting up a user account as part of approving a loan.

IBM also improved integration between Identity Manager and Directory Integrator 5.2, which is scheduled to ship next month, to increase the number of repositories that can feed user data to Identity Manager. Directory Integrator now also lets users manipulate data before passing it on, such as adding a country code to a telephone number.

IBM says next year the entire suite will be upgraded again to support Web services standards the company is developing in conjunction with Microsoft, including WS-Federation and WS-Policy. ■

Beefing up

IBM is adding enhancements to its identity management platform to improve user administration.

Product	Additions	Description
Access Manager 5.1	• Dynamic Rules Engine. • Dynamic Group Support.	Improves user management and access control.
Identity Manager 4.5	• Enhanced workflow engine.	Improves automated provisioning; can be integrated with other business workflow engines.
Directory Integrator 5.2	• Extended number of integration points.	User data can be pulled from various sources and manipulated before sharing with Identity Manager.

VoIP

can be thwarted by encrypting the voice traffic with Secure RTP.

This is key in any VoIP deployment, says Kameran Ahari, general partner in Napa Consulting Group. "True VoIP requires real-time protocol support in the context of the overall security strategy. But, the security issues are no different than some of the data applications," Ahari says.

While some might equate VoIP encryption to paranoia, it is a must for running IP voice to home users.

"At all costs, avoid going directly over the Internet" with VoIP, ThruPoint's Ortega says. If organizations want to extend access to a PBX or IP PBX to home users, encrypted VPN tunnels over a broadband link are best. ■

BackSpin Mark Gibbs

Collaboration technology: Just a lot of noise?

In a recent editorial our esteemed editorial director, John Gallant, discussed the problems of collaboration technologies and observed that "Technology is developing faster than our skills to deal with it. We're always on and always connected. But are we always better off? Are we more productive, or simply busier dealing with more messages and more distractions? Share your thoughts with me."

Oh, all right, if you insist, John, I will....

And the answer to whether we are, in general, more productive because of collaboration technologies is, I doubt it. Just consider the problems caused by instant messaging, a tool that is often a distraction and a cause of social friction.

From what my readers tell me, the instant-messaging problem is common in many organizations, where it is definitely as much a waste of time as it is a useful communications tool.

Instant messaging isn't the only problem. In many organizations e-mail has become an endless cocktail party of jokes and poor thinking. Some of this chatter is useful because it keeps people in touch and maintains relationships, but most is noise. And because the noise is mixed in with the signal it is an effort to extract information that you dare not miss.

But the problem of productivity actually has little to do with technology. Just as guns don't kill people, technology in and of itself doesn't kill productivity.

Sure, some technology is so complex, overbearing and rigid that people find it hard to use it effectively (just consider how few companies use Lotus Notes as the total enterprise information solution it was intended to be). But underlying the limitations of technology is the biggest problem of all: people.

This is because we, as human animals, are intrinsically problematic when we are collaborating. We are driven by history and biology to look for connection, to get accepted by the "tribe," to seek approval, to be wary of offense, to exercise hierarchical dominance and rivalry, and to indulge ourselves in ritualistic antagonism. And we're lazy and undisciplined. We don't take kindly to detail and concentration.

All of these drives are incredibly hard for us to put aside and very difficult to ignore in others. Worse still, our culture doesn't really frown on such traits except in the abstract. Our society is generally more concerned with style than substance and more interested in the score than how the game is played.

So mix all those human attributes with new ways of communicating and you are guaranteed to have problems. People will use these tools poorly because they don't know otherwise and their drives are usu-

ally unchecked by training or feedback.

IT groups need to make sure that communications and productivity systems are managed and their users taught how to use the tools effectively. For instant messaging and e-mail, you should run courses in their effective use. To back this up you absolutely must have acceptable-uses policies and you should ensure that the services are monitored.
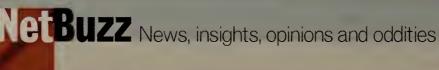
And monitoring can be very effective. By simply filtering e-mail for the Seven Dirty Words you'll be able to identify those users who are probably part of the messaging noise problem and look to manage them.

This of course leads to the problem of whether monitoring is an acceptable practice. Legal concerns aside, I'd suggest that if you have nothing to hide or be ashamed of, then you wouldn't care.

Most crucially, if corporate resources are being wasted or abused, the organization has a responsibility to fix the problem. And if that requires monitoring and correcting or even disciplining users, how bad is that? Surely that counts as a mature, commonsense solution to a serious problem?

So are we better off with these collaboration technologies? Not yet. But they are here to stay, and the sooner we start managing them effectively the sooner they will pay off.

Discipline to backspin@gibbs.com.

'NetBuzz News, insights, opinions and oddities

By Paul McNamara

They say there's a will

Executives from Iron Mountain dropped by to tout the findings of a commissioned survey that quizzed some 100 IT executives about managing business records in a world where government regulations run amok and litigation discovery can make finding that needle in a haystack painfully expensive.

Yes, it's self-serving — Iron Mountain sells records management services — but the results are interesting:

● Just about nine of every 10 execs copped to not having rolled out the tools needed to do records management right.

● More than half admitted they don't evaluate their procedures regularly.

● And, not surprisingly, 44% plan to spend more on this stuff next year; 21% substantially more.

Of course, what people say they're going to do and what they actually do often bear little resemblance, especially when time and money are involved.

For example, in those nightmarish weeks after the Sept. 11 terrorist attacks, it became an article of faith in the industry that spending on disaster recovery and business continuity was about to go on a run reminiscent of generator and battery sales just before Y2K. . . . That sense of urgency dissipated rather quickly.

"The difference with business continuity is that I still don't have to do it," says Iron Mountain President Peter Delle Donne. "In the records management piece, litigation is just overwhelming companies on a daily basis."

So, too, the weight of regulation.

"Sarbanes-Oxley is not going away," says Ken Rubin, Iron Mountain's executive vice president of marketing. "It is not a flash in the pan; it's not a Y2K type of thing. This is going to have legs, and it's just a question of when companies do [records management] right. They are already committed to it."

Uh, about that Segway recall . . .

Our 2-year-old son Grant revels in announcing that it's "dinnertime! dinnertime! dinnertime!" whenever he senses meal preparation has begun. What's especially cute is that it makes no difference whether we're nearing dinner, lunch or breakfast; it's all dinnertime to him.

While Grant works on grasping such fine distinctions, his dad is trying to decide whether now is the right time for eating crow over previous columns extolling the virtues of Segway, inventor Dean Kamen's much-ballyhooed scooter.

You might have heard that the government ordered Kamen's company to recall every Segway sold. Something about people falling off when the battery gets low.

But that recall is not the source of my angst: There probably isn't a car model on the road today that hasn't had a recall, and the fact that people fall off scooters shouldn't come as a shock.

The problem is that every Segway sold means only 6,000 of the things . . . since December 2001.

Some still insist it's way too early to write off Segway. History is chock full of important technologies that took time to take off, they say.

True enough. But 6,000 scooters?

I hear Grant calling. It's dinnertime, dinnertime, dinnertime: Pass the crow.

The Onion brings tears . . . again

Proving once more why it's the funniest site on the 'Net, here's a bite from The Onion that had my eyes watering:

48-hour Internet outage plunges nation into productivity

BOSTON — An Internet worm that disabled networks across the U.S. Monday and Tuesday temporarily thrust the nation into its most severe maelstrom of productivity since 1992.

Full "story" available at www.theonion.com.

Writing to the columnist is always productive. The address is buzz@nww.com.

Demand more from your network.
Without demanding more from your budget.

With HP ProCurve Networking solutions, you can get the secure, mobile, multiservice network you need without having to sacrifice the entire budget to get it. Built on industry standards, the HP ProCurve Switch 5300xl series is engineered to deliver wire-speed 2/3/4 functionality at the price of competitive layer 2 solutions. Our best-selling switches include HP support, software upgrades and our industry-leading lifetime warranty. And with intelligence pushed to the edge of the network, you get the control you need. All of which translates into better return on your IT investment.

HP PROCURVE SWITCH 2626

$799*

24 10/100 and 2 Gigabit ports
in a 1U form factor
Built-in Gigabit copper with fiber flexibility
Basic IP routing
Lifetime warranty**

HP PROCURVE SWITCH 5372xl

$7,129*

76.8 Gbps wire-speed switch fabric
48 Mpps throughput
Layer 3 IP routing
802.1x port-based security
HP-patented switch meshing
Lifetime warranty**

hp ®
invent

Learn more about HP ProCurve Networking solutions around security, mobility and convergence. Sign on to our Web-cast series at www.itworld.com/itwebcast/procurve.

*List price. **For as long as you own the product (available in most countries) ©2003 Hewlett-Packard Development Company, L.P.

Cut router costs in half (without compromising performance).

The NetVanta™ 3000 Series from ADTRAN™.

Dare to Compare!	NetVanta 3305	Industry-Leading Brand
Dual Network Interfaces	✔	$$$
Dual Ethernet Interfaces	✔	$$$
Stateful Inspection Firewall	✔	$$$
Command Line Interface (CLI)	✔	✔
Quality of Service (QoS)	✔	✔
VLAN Trunking	✔	✔
Virtual Private Networking (VPN)	$	$$$
Dial Backup	$	$$$
PBX Connectivity	$	$$$$$
Unlimited Telephone Support	✔	$$$
Free Maintenance Releases	✔	Not Available
Warranty	5 Year	1 Year

Uncompromising quality. Affordable price. There's no better value in access routers than the NetVanta 3000 Series from ADTRAN.

Using a NetVanta 3000 router, you can outfit a remote location with *complete T1 voice and data* communications **for 50% less than you're accustomed to paying.** Loaded with standard features, and available with very *reasonably priced options*, the NetVanta 3000 Series is everything you need in a router and more. Lower price isn't the result of cutting corners—it's the result of smart engineering. Engineering that's backed by a *100% satisfaction guarantee* from ADTRAN, including unlimited telephone technical support (before and after the sale), *free maintenance upgrades*, and a full *five-year warranty.* Try a NetVanta 3000 router today. And start getting more out of your router dollar.

Why pay more?

Test your CLI knowledge! Receive a free T-Shirt!
www.adtran.com/info/whypaymore

DARE TO COMPARE

877.767.6022 Technical Questions
877.280.8416 Where to Buy

Experts choose ADTRAN™   **ADTRAN**®